arXiv:1701.02127v1 [cs.CV] 9 Jan 2017

Discrete approximations of affine Gaussian receptive fields*

Tony Lindeberg

Computational Brain Science Lab
Department of Computational Science and Technology
School of Computer Science and Communication
KTH Royal Institute of Technology
SE-100 44 Stockholm, Sweden

Abstract

This paper presents a theory for discretizing the affine Gaussian scale-space concept so that scale-space properties hold also for the discrete implementation.

Two ways of discretizing spatial smoothing with affine Gaussian kernels are presented: (i) by solving semi-discretized affine diffusion equation as derived by necessity from the requirement of a semi-group structure over a continuum of scale parameters as parameterized by a family of spatial covariance matrices and obeying non-creation of new structures from any finer to any coarser scale as formalized by the requirement of non-enhancement of local extrema and (ii) a set of parameterized 3×3-kernels as derived from an additional discretization of the above theory along the scale direction and with the parameters of the kernels having a direct interpretation in terms of the covariance matrix of the composed discrete smoothing operation.

We show how convolutions with the first family of kernels can be implemented in terms of a closed form expression for the Fourier transform and analyse how a remaining degree of freedom in the theory can be explored to ensure a positive discretization and optionally also achieve higher-order discrete approximation of the angular dependency of the shapes of the affine Gaussian kernels.

We do also show how discrete directional derivative approximations can be efficiently implemented to approximate affine Gaussian derivatives as constituting a canonical model for receptive fields over a purely spatial image domain and with close relations to receptive fields in biological vision.

Keywords: scale space, scale, affine, receptive field, Gaussian kernel, discrete, spatial, spatio-chromatic, double-opponent, feature detection, computer vision.

*The support from the Swedish Research Council (contract 2014-4083) is gratefully acknowledged.

1 Introduction

A basic fact when interpreting image information is that pointwise measurement of the image intensity $f(x, y)$ at a single image point does usually not carry any relevant information, since it is dependent on external and unknown illumination. The essential information is instead mediated by the relations between the image intensities at neighbouring points. To handle this issue, the notion of receptive fields has been developed in both biological vision and computer vision by performing local image measurements over neighbourhoods of different spatial extent (Hubel and Wiesel [25, 26]; DeAngelis *et al.* [14, 13]; Conway and Livingstone [7]; Johnson *et al.* [28]; Shapley and Hawken *et al.* [60]; Koenderink and van Doorn [30, 31]; Young *et al.* [75, 76]; Lindeberg [39, 45]; Schiele and Crowley [58]; Lowe [51]; Dalal and Triggs [12]; Serre *et al.* [59]; Bay *et al.* [3]; Burghouts and Geusebroek [5]; van de Sande *et al.* [57]; Tola *et al.* [64]; Linde and Lindeberg [36]; Larsen *et al.* [33]; Krizhevsky *et al.* [32]; Simonyan and Zisserman [61]; Szegedy *et al.* [63]; He *et al.* [23]).

A fundamental problem in this context concerns what shapes of receptive field profiles are meaningful. Would any shape of the receptive field do? This problem has been addressed from an axiomatical viewpoint in the area of scale-space theory, by deriving families of receptive fields that obey scale-space axioms that reflect structural properties of the world. Over an isotropic spatial domain, arguments and axiomatic derivations by several researchers have shown that the Gaussian kernel and Gaussian derivatives derived from it constitutes a canonical family of kernels to model spatial receptive fields (Iijima [27]; Witkin [74]; Koenderink [29]; Koenderink and van Doorn [30, 31]; Lindeberg [39, 40, 43, 44, 46]; Florack [19]; ter Haar Romeny [21]; Weickert *et al.* [73]). Such Gaussian derivatives can in turn be used as primitives for expressing a large class of visual modules in computer vision, including feature detection, feature classification, surface shape, image matching and object recognition.

The underlying assumption about spatial isotropy that underlies the part of the scale-space theory that leads to receptive fields based on rotational symmetric Gaussian kernels is, however, neither necessary nor always desirable. For images of objects that are subject to variations in the viewing direction, the perspective mapping leads to non-isotropic perspective transformations that are not within the group of image transformations covered by rotationally symmetric Gaussians. If we locally at every point approximate the non-linear perspective mapping by its derivative, the effect of the perspective mapping can for locally smooth surfaces to first order of approximations be approximated by local affine image deformations. If we would like to base a vision system on a receptive field model that is closed under affine transformations, we should therefore replace the rotationally symmetric Gaussian kernels by affine Gaussian kernels (Lindeberg [39]; Lindeberg and Gårding [49, 50]). Interestingly, the family of affine Gaussian kernel and affine Gaussian derivatives can also be uniquely derived from basic scale-space axioms that reflect structural properties of the world in combination with a requirement about internal consistency between image representations at multiple spatial scales (Lindeberg [44]).

The subject of this article is to address the problem of how to discretize receptive fields based on affine Gaussian derivatives while preserving scale-space properties also in a discrete sense. This work constitutes a continuation of a previously developed scale-space theory for discrete signals and images (Lindeberg [37, 39, 38]), which was conceptually extended to a discrete affine scale-space in a conference publication (Lindeberg [41]), while leaving out several details that are needed when implementing the theory in practice. The work is also closely related to other work on discretizing scale-space operations as done by Deriche [15], Weickert *et al.* [70, 72], van Vliet *et*

al. [68], Wang [69], Florack [18], Almansa and Lindeberg [1], Lim and Stiehl [35], Geusebroek *et al.* [20], Farid and Simoncelli [17], Bouma *et al.* [4], Dollár *et al.* [16], Tschirsich and Kuijper [65], Slavík and Stehlík [62], Lindeberg [47] and others.

It will be shown that by an implementation in terms of an explicit closed-form expression for the discrete Fourier transform, it is possibly to exactly compute convolutions with discrete affine kernels up to the numerical errors in a fast Fourier transform and that this can be achieved over a continuum over the multi-dimensional spatial scale parameter as parameterized by a family of spatial covariance matrices. It will also be shown that by a further discretization of this theory over the scale direction, it is possible to derive a model of compact 3×3-kernels that can be interpreted as discretizations of affine Gaussian kernels. Such kernels have a close relationship to recent developments in deep learning, where deep networks such as VGG-Net (Simonyan and Zisserman [61]) and ResNet (He *et al.* [23]) are specifically based on automated learning of compact 3×3 kernels that are repeatedly applied to the image data. With the theory presented in this paper, we open up for relating such kernels to scale-space operations, with the possibilities of building closer links between the empirical developments in deep learning and axiomatically derived scale-space theory with its in turn close relations to receptive fields in biological vision.

2 Affine Gaussian receptive fields over a continuous image domain

Starting from assumptions about (i) linearity, (ii) shift invariance, (iii) a semi-group structure over spatial scales and (iv) a requirement of not introducing new structures from finer to coarser scales in terms of non-enhancement of local extrema, which means that the value at a local maximum must not increase from finer to coarser scales and correspondingly the value at a local minimum must not decrease, it can be shown that affine Gaussian kernels

$$g(x;\; \Sigma_s) = \frac{1}{2\pi\sqrt{\det \Sigma_s}}\, e^{-x^T \Sigma_s^{-1} x/2}, \tag{1}$$

and their directional derivatives constitute a canonical model for spatial receptive fields over a two-dimensional image domain (Lindeberg [44]). For the specific parameterization of the spatial covariance matrix $\Sigma_s = s\Sigma$, convolution these kernels obeys the affine diffusion equation

$$\partial_s L = \frac{1}{2} \nabla^T \left(\Sigma\, \nabla L\right). \tag{2}$$

with initial condition $L(x, y;\; 0) = f(x, y)$. Convolution with affine Gaussian kernels also obeys the cascade smoothing property

$$g(\cdot, \cdot;\; \Sigma_1) * L(\cdot, \cdot;\; \Sigma_2) = L(\cdot, \cdot;\; \Sigma_1 + \Sigma_2). \tag{3}$$

To parameterize the affine Gaussian kernels, let us in the two-dimensional case consider the covariance matrix determined by two eigenvalues λ_1, λ_2 and one orientation α. Then, the covariance matrix can be written

$$\Sigma' = \begin{pmatrix} \lambda_1 \cos^2 \alpha + \lambda_2 \sin^2 \alpha & (\lambda_1 - \lambda_2) \cos \alpha\, \sin \alpha \\ (\lambda_1 - \lambda_2) \cos \alpha\, \sin \alpha & \lambda_1 \sin^2 \alpha + \lambda_2 \cos^2 \alpha \end{pmatrix}. \tag{4}$$

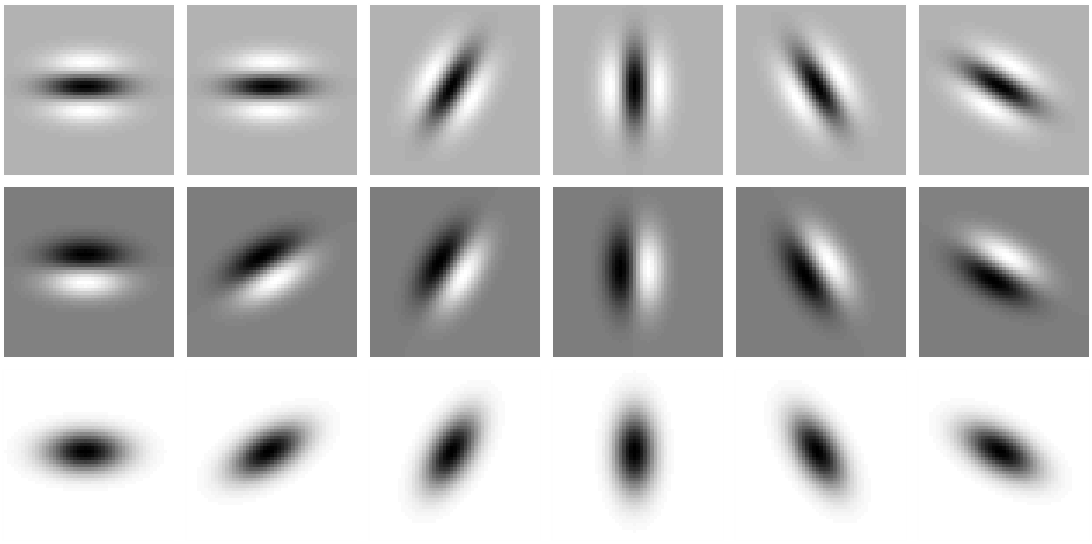

Figure 1: Examples of affine Gaussian kernels $g(x, y;\ \Sigma)$ and their directional derivatives $\partial_{\perp\varphi} g(x, y;\ \Sigma)$ and $\partial_{\perp\varphi\perp\varphi} g(x, y;\ \Sigma)$ up to order two in the two-dimensional case, here for $\lambda_1 = 64$, $\lambda_2 = 16$ and $\alpha = 0, \pi/6, \pi/3, \pi/2, 2\pi/3, 5\pi/6$. (Horizontal axis: $x \in [-24, 24]$. Vertical axis: $y \in [-24, 24]$.)

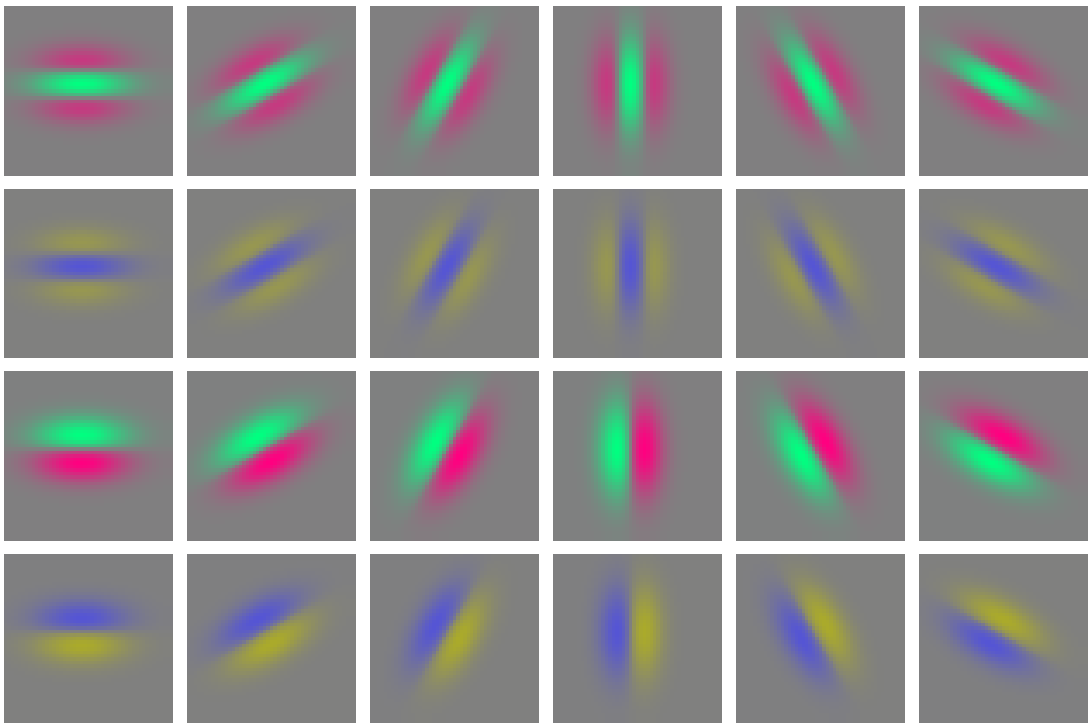

Figure 2: Examples of affine Gaussian colour-opponent directional derivatives according to (8) and (7) up to order two in the two-dimensional case, here for $\lambda_1 = 64$, $\lambda_2 = 16$ and $\alpha = 0, \pi/6, \pi/3, \pi/2, 2\pi/3, 5\pi/6$. (Horizontal axis: $x \in [-24, 24]$. Vertical axis: $y \in [-24, 24]$.)

Correspondingly, we can for two orthogonal directions φ and $\perp\varphi$ aligned to the eigendirections of the spatial covariance matrix parameterize the directional derivative operators in these directions according to

$$\partial_{\varphi} = \cos\varphi\, \partial_x + \sin\varphi\, \partial_y \qquad \partial_{\perp\varphi} = -\sin\varphi\, \partial_x + \cos\varphi\, \partial_y. \tag{5}$$

This leads to the following canonical model for affine Gaussian receptive fields over a spatial domain (Lindeberg [44])

$$g_{\varphi^m \perp\varphi^n}(x, y;\ \Sigma_s) = \partial^m_{\varphi} \partial^n_{\perp\varphi}\, g(x;\ \Sigma_s), \tag{6}$$

preferably with the angle φ in (5) and (6) chosen equal to the angle α in (4). Figure 1 shows a few examples of affine Gaussian kernels and their directional derivatives up to order two for one combination of the two eigenvalues determined such that the ratio between the scale parameters in the two eigendirections of the covariance matrix is equal to two, with the scale parameters $\sigma_1 = \sqrt{\lambda_1}$ and $\sigma_2 = \sqrt{\lambda_2}$ expressed in terms of dimension [length].

Consider next the colour channels of a colour-opponent space defined from an RGB colour representation (Hall *et al.* [22])

$$\begin{pmatrix} f \\ u \\ v \end{pmatrix} = \begin{pmatrix} f \\ c^{(1)} \\ c^{(2)} \end{pmatrix} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & -\frac{1}{2} & 0 \\ \frac{1}{2} & \frac{1}{2} & -1 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}, \tag{7}$$

where yellow is approximated by the average of the R and G channels and f can be defined as a channel of pure intensity information. Then, *affine Gaussian colour-opponent receptive fields* $(U, V)^T = (C^{(1)}, C^{(2)})^T$ can be defined by applying affine Gaussian receptive fields of the form (6) to the colour-opponent channels $(c^{(1)}, c^{(2)})$:

$$U = C^{(1)}(\cdot, \cdot;\ \Sigma_s) = g_{\varphi^m \perp\varphi^n}(\cdot, \cdot;\ \Sigma_s) * c^{(1)}(\cdot), \tag{8}$$

$$V = C^{(2)}(\cdot, \cdot;\ \Sigma_s) = g_{\varphi^m \perp\varphi^n}(\cdot, \cdot;\ \Sigma_s) * c^{(2)}(\cdot). \tag{9}$$

Figure 2 shows examples of such affine Gaussian spatio-chromatic receptive fields up to order two over red/green and yellow/blue colour-opponent space.

More generally, one should consider affine receptive fields as being present for all positive definite covariance matrices Σ as parameterized by their eigenvalues $\lambda_1 > 0$ and $\lambda_2 > 0$ and for all directions α. This model is proposed in Lindeberg [44, Section 6] [45, Section 6.3]) as a model for the spatial component of simple cells in the primary visual cortex (V1) and in good agreement with neural cell recordings by DeAngelis *et al.* [14, 13] and Johnson *et al.* [28]. Figure 3 shows an example of the spatial dependency of a simple cell that can be well modelled by a first-order affine Gaussian derivative over image intensities. Figure 4 shows corresponding results for a color-opponent receptive field of a simple cell in V1 that can be modelled as a first-order affine Gaussian spatio-chromatic derivative over an R-G colour-opponent channel.

A theoretically very attractive property of the affine Gaussian model for spatial receptive fields is that this family of receptive fields is closed under affine transformations of the spatial domain. If we consider two images f_L and f_R that with vector notation for the image coordinates $\xi = (\xi_1, \xi_2)^T$ and $\eta = (\eta_1, \eta_2)^T$ are related by an affine image deformation

$$f_L(\xi) = f_R(\eta) \quad \text{where} \quad \eta = A\, \xi + b \tag{10}$$

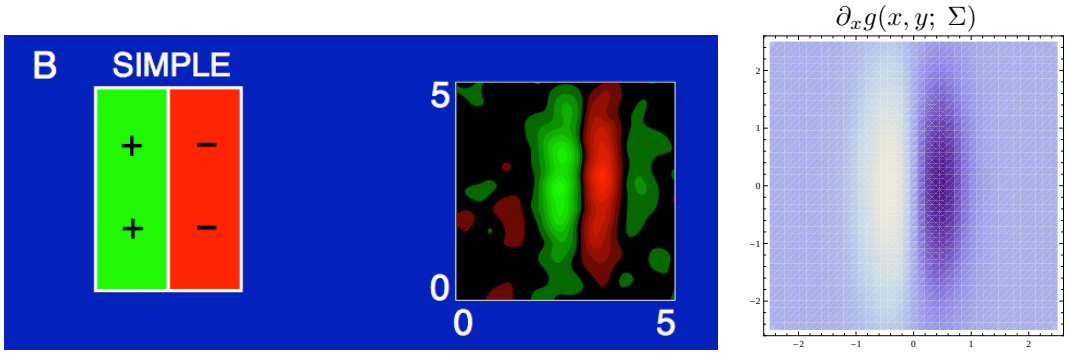

Figure 3: Example of a receptive field profile over the spatial domain in the primary visual cortex (V1) as reported by DeAngelis *et al.* [14, 13]. (middle) Receptive field profile of a simple cell over image intensities as reconstructed from cell recordings, with positive weights represented as green and negative weights by red. (left) Stylized simplification of the receptive field shape. (right) Idealized model of the receptive field from a first-order directional derivative of an affine Gaussian kernel $\partial_x g(x, y;\ \Sigma) = \partial_x g(x, y;\ \lambda_x, \lambda_y) = -\frac{x}{\lambda_x} 1/(2\pi\sqrt{\lambda_x \lambda_y}) \exp(-x^2/2m\lambda_x - y^2/2\lambda_y)$, here with $\lambda_x = 0.2$ and $\lambda_y = 2$ in units of degrees of visual angle, and with positive weights with respect to image intensities represented by white and negative values by violet.

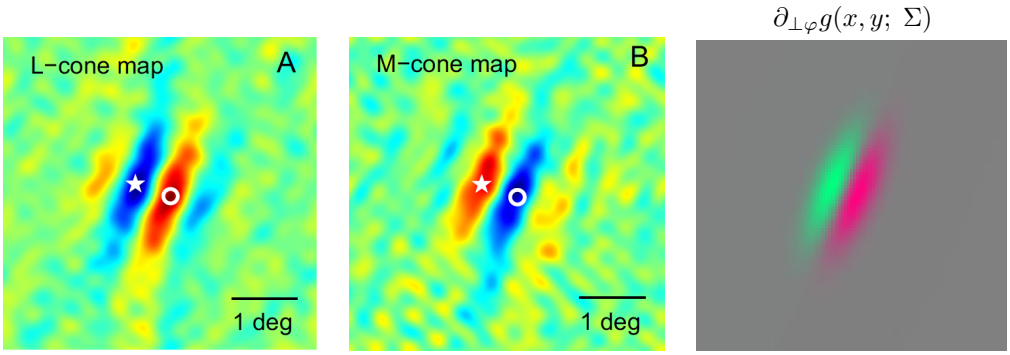

Figure 4: Example of a colour-opponent receptive field profile over the spatial domain for a double-opponent simple cell in the primary visual cortex (V1) as measured by Johnson *et al.* [28]. (left) Responses to L-cones corresponding to long wavelength red cones, with positive weights represented by red and negative weights by blue. (middle) Responses to M-cones corresponding to medium wavelength green cones, with positive weights represented by red and negative weights by blue. (right) Idealized model of the receptive field from a first-order directional derivative of an affine Gaussian kernel $\partial_{\perp\varphi} g(x, y;\ \Sigma)$ according to (6), (5), (1) and (4) for $\sigma_1 = \sqrt{\lambda_1} = 0.6$, $\sigma_2 = \sqrt{\lambda_2} = 0.2$ in units of degrees of visual angle, $\alpha = 67$ degrees and with positive weights for the red-green colour-opponent channel U according to (8) and (7) represented by red and negative values by green.

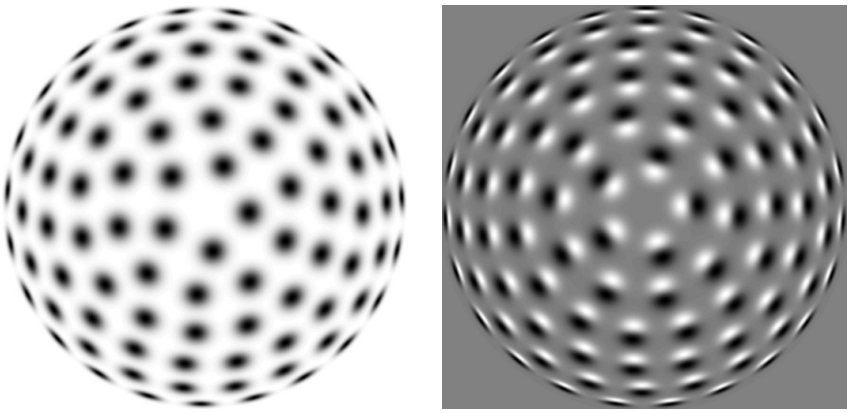

Figure 5: Distributions of affine Gaussian receptive fields corresponding to a uniform distribution on a hemisphere regarding (left) zero-order smoothing kernels and (right) first-order derivatives.

define corresponding scale-space representations according to

$$L(\cdot;\, \Sigma_L) = g(\cdot;\, \Sigma_L) * f_L(\cdot), \quad R(\cdot;\, \Sigma_R) = g(\cdot;\, \Sigma_R) * f_R(\cdot), \tag{11}$$

then these scale-space representations will be related according to (Lindeberg [39]; Lindeberg and Garding [50])

$$L(x;\, \Sigma_L) = R(y;\, \Sigma_R) \quad \text{where} \quad \Sigma_R = A\, \Sigma_L\, A^T \quad \text{and} \quad y = A\, x + b \tag{12}$$

where $x = (x_1, x_2)^T$ and $y = (y_1, y_2)^T$. In other words, given that an image f_L is affine transformed into an image f_R it will always be possible to find a transformation between the spatial covariance matrices Σ_L and Σ_R in the two domains that makes it possible to match the corresponding derived internal representations $L(\cdot;\, s_L)$ and $R(\cdot;\, s_R)$. If we in turn locally approximate the non-linear perspective image deformation between two images of the same scene by local affine deformations, this means that it will be possible to match the receptive field responses computed from different views of the same scene.

This idea was originally proposed for reducing the shape distortions that arise when computing estimates of local surface orientation by shape-from-texture and shape-from-disparity-gradients (Lindeberg and Garding [49, 50]) and has later been explored for image-based matching and recognition (Baumberg [2]; Mikolajczyk and Schmid [52]; Tuytelaars and van Gool [66]; Mikolajczyk *et al.* [53]; Lazebnik *et al.* [34]; Rothganger *et al.* [55]; Tuytelaars and Mikolajczyk [67]). Similar ideas of affine invariance as underlying these approaches based on affine shape adaptation have also been used for designing smoothing methods with a larger amount of smoothing along local image structures than across them (Weickert [71]; Almansa and Lindeberg [1]) and for constructing affine SIFT descriptors (Morel and Yu [54]; Sadek *et al.* [56]).

When implementing such affine covariance in practice, there are two main approaches to follow, either by adapting the shape of the affine Gaussian kernel as proposed in the affine shape adaptation methodology proposed in (Lindeberg and Garding [49, 50]) or by representing all affine covariances matrices alternatively some sparse sampling of the resulting family of receptive fields. Figure 5 shows the resulting distributions of affine Gaussian receptive fields of different orientations and degrees of orientation as they arise from local linearizations of a perspective projection model if we assume that the set of surface directions in the world is on average uniformly

distributed in the world and if the distributions of the local surface patterns on these object surfaces are in turn without dominant directional bias and uncoupled to the orientations of the local surface patches. In our idealized model of receptive fields, all these receptive fields can be thought of as being present at every position in image space, and corresponding to a uniform distribution on a hemisphere.

3 Discretizing affine Gaussian receptive fields

When implementing affine Gaussian receptive fields computationally, there are different approaches to follow. The most straightforward approach consists of sampling the affine Gaussian kernel or its derivatives at the grid points of the image domain or alternatively integrating these kernels over the support region of each image pixel. It is, however, known that sampling the Gaussian kernel is not the best way of implementing the regular scale-space concept based on rotationally symmetric Gaussian kernels, where local integration of the scale-space kernel is a better choice that remedies some of the artifacts (Lindeberg [37]). More seriously with respect to computational efficiency, the resulting kernels will in general be inseparable and will therefore for truncated filter size of M along each dimension require M^2 operations as opposed to $2M$ operations for separable filtering. Implementing the affine Gaussian kernels in the Fourier domain is then an alternative option, but besides the technical need for additional extensions when handling image sizes that do not match well with even powers of 2, there are then also discretization issues to consider.

In this and following sections, we will develop a spatial discretization approach based on discrete scale-space theory that has been specifically designed to preserve scale-space properties to the discretization (Lindeberg [41, 42]), and then discretize the corresponding semi-discrete scale-space over a continuum of spatial scale levels further to compact 3×3 kernels that are to be applied repeatedly and will lead to discrete scale-space kernels whose covariance matrices are exactly equal to the covariance matrices that would be obtained from the corresponding continuous theory.

An underlying motivation is to replace the training of the earliest layers in deep networks by theoretically derived scale-space filters and here specifically replace the repeated application of trained 3×3 kernels in VGG-Net (Simonyan and Zisserman [61]) and ResNet (He *et al.* [23]) by kernels that obey the algebra of affine Gaussian receptive fields to enable affine covariance of the deep network to in turn enable provably affine invariant recognition methods based on deep learning.

3.1 Continuous scale parameter scale space over a discrete image domain

In (Lindeberg [41, 42]) different generalizations are presented of the discrete scale-space theory originally proposed in (Lindeberg [37, 39]) for an isotropic spatial image domain and leading to rotationally symmetric scale-space kernels to non-isotropic spatial domains where the scale-space kernels are no longer required to be rotationally symmetric as well as to spatio-temporal image domains where the temporal dimension is treated in a conceptually different way than the spatial dimensions.

For signals defined over a D-dimensional discrete domain, it can be shown that non-enhancement of local extrema implies that the scale-space family $L \colon \mathbb{Z}^D \times \mathbb{R}^+ \to \mathbb{R}$ of any discrete signal $f \colon \mathbb{Z}^D \to \mathbb{R}$ must satisfy the semi-discrete differential equation

$$(\partial_s L)(x;\ s) = (\mathcal{A} L)(x;\ s) = \sum_{\xi \in \mathbb{Z}^D} a_\xi L(x - \xi;\ s), \tag{13}$$

for some *infinitesimal scale-space generator* $\mathcal{A}$ characterized by

- the *locality* condition $a_\xi = 0$ if $|\xi|_\infty > 1$,
- the *positivity* constraint $a_\xi \geq 0$ if $\xi \neq 0$, and
- the *zero sum* condition $\sum_{\xi \in \mathbb{Z}^D} a_\xi = 0$.

This result follows from similar arguments as the proof of Theorem 4.10 in (Lindeberg [39]), if the spatial symmetry requirements are relaxed (see appendix A for a formal proof). If we additionally require a spatial image domain to be mirror symmetric around the origin, then we should additionally require

- the *symmetry* condition $a_{-\xi} = a_\xi$.

To investigate the solutions of this equation, a general approach that we shall follow in this paper consists of computing the generating function

$$\varphi(z;\ s) = \sum_{n \in \mathbb{Z}^D} c_n z^n \tag{14}$$

of the set of filter coefficients c_n in the filter for computing the scale-space representation L from the input signal f

$$L(x;\ s) = \sum_{n \in \mathbb{Z}^D} c_n\, f(x - n). \tag{15}$$

By formally transforming (13) into generating functions

$$\partial_s \left(\varphi_C(z;\ s)\right) = \left(\sum_{\xi \in \mathbb{Z}^D} a_\xi z^\xi \right) \varphi_C(z;\ s) \tag{16}$$

where $a_\xi z^\xi$ should be interpreted as multi-index notation of $a_{\xi_1,\ldots,\xi_D} z_1^{\xi_1} \ldots z_D^{\xi_D}$, we obtain

$$\varphi(z;\ s) = e^{s \sum_{\xi \in \mathbb{Z}^D} a_\xi z^\xi}. \tag{17}$$

In compact operator notion, the solution of (13) at scale s can equivalently be written

$$L = e^{s\,\mathcal{A}} f \tag{18}$$

and this expression can be regarded as a general parameterization of scale-space kernels on discrete image domains. Alternatively, we obtain the Fourier transform of the kernel by substituting $z = e^{-i\omega}$ into the generating function with the multi-index interpretation $z^\xi = z_1^{\xi_1} \ldots z_D^{\xi_D} = e^{-i\omega_1 \xi_1} \ldots e^{-i\omega_D \xi_D} = e^{-i\omega^T \xi}$, which gives

$$\psi(\omega;\ s) = \sum_{n \in \mathbb{Z}^D} c_n e^{-i\omega^T n} = e^{s \sum_{\xi \in \mathbb{Z}^D} a_\xi e^{-i\omega^T \xi}}. \tag{19}$$

A main subject of this article is to interpret this result over a two-dimensional image domain and develop a theory for discrete affine scale space.

The methodology we shall follow is to reparametrize the filter class in terms of the following basic difference operators, here over the dimension x and analogously for the dimension y

$$(\delta_x f)(x) = \left(f(x+1) - f(x-1)\right)/2, \tag{20}$$

$$(\delta_{xx} f)(x) = f(x+1) - 2f(x) + f(x-1), \tag{21}$$

which will be used as a basis for expressing the degrees of freedom in the coefficients a_ξ.

4 Discrete two-dimensional affine Gaussian scale space

For a two-dimensional spatial domain, the discrete counterpart of the affine Gaussian scale-space in (1) and (2) is obtained if we require the discrete filter kernels in (13) to be mirror symmetric through the origin, *i.e.*, $a_{i,j} = a_{-i,-j}$. Then, the computational molecule of the infinitesimal operator $\mathcal{A}$ can be written as

$$\mathcal{A} = \begin{pmatrix} a_{-1,1} & a_{0,1} & a_{1,1} \\ a_{-1,0} & a_{0,0} & a_{1,0} \\ a_{-1,-1} & a_{0,-1} & a_{1,-1} \end{pmatrix} = \begin{pmatrix} D & C & B \\ A & -E & A \\ B & C & D \end{pmatrix} \tag{22}$$

for some $A, B, C, D \geq 0$ and $E = A+B+C+D$. In terms of the previously mentioned difference operators, and with $\delta_{xy} = \delta_x \, \delta_y$ and $\delta_{xxyy} = \delta_{xx} \, \delta_{yy}$, this computational molecule can with

$$A = \frac{C_{xx}}{2} - \frac{C_{xxyy}}{2}, \tag{23}$$

$$B = \frac{C_{xy}}{4} + \frac{C_{xxyy}}{4}, \tag{24}$$

$$C = \frac{C_{yy}}{2} - \frac{C_{xxyy}}{2}, \tag{25}$$

$$D = -\frac{C_{xy}}{4} + \frac{C_{xxyy}}{4}, \tag{26}$$

be reparameterized as

$$\mathcal{A} = \frac{1}{2} \begin{pmatrix} -C_{xy}/2 & C_{yy} & C_{xy}/2 \\ C_{xx} & -2(C_{xx} + C_{yy}) & C_{xx} \\ C_{xy}/2 & C_{yy} & -C_{xy}/2 \end{pmatrix} + \frac{C_{xxyy}}{4} \begin{pmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{pmatrix}. \tag{27}$$

4.1 Semi-discrete affine diffusion equation

With the parameterization in terms of C_{xx}, C_{xy}, C_{yy} and C_{xxyy}, the corresponding *discrete affine Gaussian scale-space* is given as the solution of the semi-discrete differential equation

$$\partial_s L = \frac{1}{2} \left(C_{xx} \, \delta_{xx} L + 2 C_{xy} \, \delta_{xy} L + C_{yy} \, \delta_{yy} L \right) + \frac{C_{xxyy}}{4} \, \delta_{xxyy} L, \tag{28}$$

which in turn can be interpreted as a second-order discretization of the diffusion equation associated with the continuous affine Gaussian scale-space

$$\partial_s L = \tfrac{1}{2} (C_{xx} \, L_{xx} + 2 C_{xy} \, L_{xy} + C_{yy} \, L_{yy}), \tag{29}$$

where $C_{xx}, C_{yy} > 0$ and $C_{xx} C_{yy} - C_{xy}^2 > 0$ are necessary conditions for the infinitesimal operator $\mathcal{A}$ to be positive definite.

4.2 Generating function and Fourier transform

With the variable for the generating function over the x-dimension denoted by z and the variable for the generating function over the y-dimension denoted by w, the generating function (17) as parametrized by the coefficients C_{xx}, C_{xy}, C_{yy} and C_{xxyy} becomes

$$\begin{aligned} \varphi(z, w) = \exp(& C_{xx} \, (z - 2 + z^{-1})/2 + C_{yy} \, (w - 2 + w^{-1})/2 + \\ & C_{xy} \, (z - z^{-1})(w - w^{-1})/4 + C_{xxyy} \, (z - 2 + z^{-1}) \, (w - 2 + w^{-1})/4). \end{aligned} \tag{30}$$

From well-known properties of the generating function, the corresponding discrete kernels have spatial mean vector

$$m = \begin{pmatrix} m_x \\ m_y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{31}$$

and spatial covariance matrix

$$\Sigma = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy} & \Sigma_{yy} \end{pmatrix} = \begin{pmatrix} C_{xx} & C_{xy} \\ C_{xy} & C_{yy} \end{pmatrix}. \tag{32}$$

Thus, by this definition of discrete affine Gaussian scale space, provided that the filter parameters C_{xx}, C_{xy}, C_{yy} and C_{xxyy} are chosen such that the semi-discrete affine diffusion equation (28) corresponds to a non-negative discretization of the continuous affine diffusion equation (29), the mean vectors and the covariances of the corresponding discrete affine Gaussian kernels are *exactly equal* to the mean values and covariance matrices of the corresponding continuous affine Gaussian kernels.

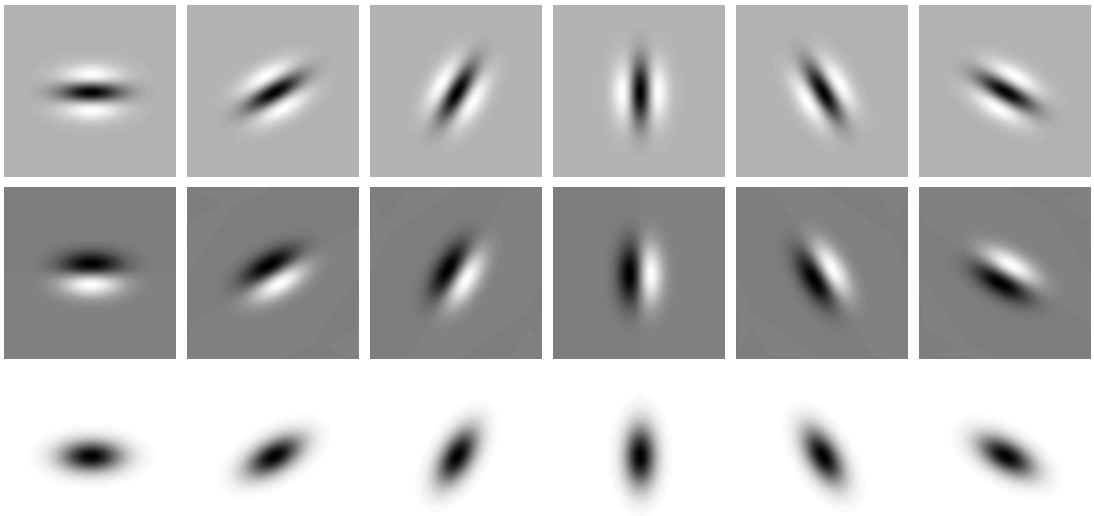

Figure 6: Examples of discrete affine Gaussian kernels $h(x, y;\, \Sigma)$ and their equivalent directional derivative approximation kernels $\delta_{\perp\varphi} h(x, y;\, \Sigma)$ and $\delta_{\perp\varphi\perp\varphi} h(x, y;\, \Sigma)$ up to order two in the two-dimensional case, here as generated from the explicit expression of the Fourier transform (34) for $\lambda_1 = 64$, $\lambda_2 = 16$ and $\alpha = 0, \pi/6, \pi/3, \pi/2, 2\pi/3, 5\pi/6$ and with $C_{xxyy} = |C_{xy}|$ according to (52). (Kernel size: 64×64 pixels.)

From the explicit expression for the Fourier transform, we can in turn obtain the continuous Fourier transform over an infinite discrete spatial domain by letting $(z, w) = (e^{iu}, e^{iv})$, which gives

$$\begin{aligned} \psi(u, v) &= \varphi(e^{iu}, e^{iv}) \\ &= \exp(-C_{xx}\,(1 - \cos u) - C_{yy}\,(1 - \cos v) + C_{xy} \sin u \sin v \\ &\qquad + C_{xxyy}\,(1 - \cos u)\,(1 - \cos v)). \end{aligned} \tag{33}$$

The corresponding discrete Fourier transform over a finite image of size $M \times N$ can then be obtained by setting $u = 2m\pi/M$ and $v = 2n\pi/N$ for $m = 0 \ldots M - 1$ and $n = 0 \ldots N - 1$, which gives

$$\begin{aligned} \psi(m, n) &= \varphi(e^{\frac{i2m\pi}{M}}, e^{\frac{i2n\pi}{N}}) \\ &= \exp(-C_{xx}\,(1 - \cos \tfrac{2m\pi}{M}) - C_{yy}\,(1 - \cos \tfrac{2n\pi}{N}) + C_{xy} \sin \tfrac{2m\pi}{M} \sin \tfrac{2n\pi}{N} \\ &\qquad + C_{xxyy}\,(1 - \cos \tfrac{2m\pi}{M})\,(1 - \cos \tfrac{2n\pi}{N})). \end{aligned} \tag{34}$$

Besides numerical errors in an FFT algorithm, we can thereby from these explicit expressions for the Fourier transform of the discrete affine Gaussian kernels generate discrete affine Gaussian kernels and compute convolutions with them with spatial covariance matrices that are exactly equal to the corresponding spatial covariance matrices of the continuous theory. What remains is to determine the free parameter C_{xxyy} in the spatial discretization as function of the parameters C_{xx}, C_{xy} and C_{yy} that determine the spatial covariance matrix of the discrete affine Gaussian kernel.

4.3 Free parameter and positivity constraint

There is one free parameter $C_{xxyy} \geq 0$ in (28), which controls the addition of a discretization of the mixed fourth-order derivative L_{xxyy}. From (23)–(26) and the positivity condition $A, B, C, D \geq 0$ it follows that this parameter must satisfy

$$|C_{xy}| \leq C_{xxyy} \leq \min(C_{xx}, C_{yy}). \tag{35}$$

In practice, the feasibility condition $|C_{xy}| \leq \min(C_{xx}, C_{yy})$ arising from this positivity constraint is always more restrictive than the condition $|C_{xy}| \leq \sqrt{C_{xx}\, C_{yy}}$ for the infinitesimal operator in (29) to be positive semi-definite. This implies that highly eccentric affine Gaussian kernels cannot be represented by non-negative discrete scale-space kernels on a square discrete grid, unless the orientation of the filter is approximately aligned to the coordinate directions.

To analyse for which positive definite covariance matrices the positivity condition may be violated, let us reparameterize the covariance matrix in terms of its eigenvalues $\lambda_{1,2} > 0$ as well as the orientation α of the eigenvectors

$$C_{xx} = \lambda_1 \cos^2 \alpha + \lambda_2 \sin^2 \alpha, \tag{36}$$

$$C_{xy} = (\lambda_1 - \lambda_2) \cos \alpha \, \sin \alpha, \tag{37}$$

$$C_{yy} = \lambda_1 \sin^2 \alpha + \lambda_2 \cos^2 \alpha. \tag{38}$$

Then, the positivity requirement $|C_{xy}| \leq \min(C_{xx}, C_{yy})$ assumes the form

$$|\lambda_1 - \lambda_2| \, |\cos \alpha \, \sin \alpha| \leq \min(\lambda_1 \cos^2 \alpha + \lambda_2 \sin^2 \alpha, \lambda_1 \sin^2 \alpha + \lambda_2 \cos^2 \alpha), \tag{39}$$

which can be rewritten as

$$|\lambda_1 - \lambda_2| \, |\sin 2\alpha| \leq \min(\lambda_1 + \lambda_2 + (\lambda_1 - \lambda_2) \cos 2\alpha, \lambda_1 + \lambda_2 - (\lambda_1 - \lambda_2) \cos 2\alpha) \tag{40}$$

or equivalently

$$|\lambda_1 - \lambda_2| \, (|\cos 2\alpha| + |\sin 2\alpha|) \leq \lambda_1 + \lambda_2. \tag{41}$$

The "worst case orientation" is given by $\alpha = \frac{\pi}{8}$, and violations of the positivity requirement start to occur when λ_1 and λ_2 obey the relation

$$\sqrt{2} \, |\lambda_1 - \lambda_2| = \lambda_1 + \lambda_2, \tag{42}$$

corresponding to a ratio of the eigenvalues equal to

$$\epsilon_{pos-bound} = \frac{\lambda_{max}}{\lambda_{min}} = 3 + 2\sqrt{2} \approx 5.8. \tag{43}$$

As long as the eccentricity of the affine Gaussian kernels is lower than this value, we can, however, represent the affine Gaussian scale-space kernels by a non-negative discretization.

A similar bound on the positivity requirement has been derived by Weickert [70] in the context of non-negative discretizations of non-linear diffusion equations.

If we consider the task of representing the distribution of affine Gaussian kernels as parameterized by spatial covariance matrices on the hemisphere as illustrated in figure 5, then the bound (43) on the eccentricity ϵ as obtained from a positivity constraint corresponds to an angle on the hemisphere relative to the north pole of

$$\theta = \arccos \frac{1}{\sqrt{\epsilon_{pos-bound}}} = 65.5 \text{ degrees}. \tag{44}$$

4.4 Behaviour for low frequencies

By transforming the Fourier transform (34) of the discrete affine Gaussian kernel to polar coordinates $(u, v) = \omega(\cos\beta, \sin\beta)$, and performing a Taylor expansion for low angular frequencies ω

$$\begin{aligned}\psi(\omega\cos\beta, \omega\sin\beta) = &- \frac{1}{2}\left(C_{xx}\cos^2\beta + 2C_{xy}\cos\beta\sin\beta + C_{yy}\sin^2\beta\right)\omega^2 \\ &+ \frac{1}{24}\left(C_{xx}\cos^4\beta + 4C_{xy}\cos^3\beta\sin\beta + 6C_{xxyy}\cos^2\beta\sin^2\beta\right. \\ &\left. + 4C_{xy}\cos\beta\sin^3\beta + C_{yy}\sin^4\beta\right)\omega^4 + \mathcal{O}(\omega^6),\end{aligned} \tag{45}$$

we can with $C_{xxyy} = (C_{xx} + C_{yy})/6 + \rho/3$ rewrite this expression as

$$\begin{aligned}\psi(\omega\cos\beta, \omega\sin\beta) = &- \frac{1}{2}\left(C_{xx}\cos^2\beta + 2C_{xy}\cos\beta\sin\beta + C_{yy}\sin^2\beta\right)\omega^2 \\ &+ \frac{1}{24}\left(C_{xx}\cos^2\beta + 2C_{xy}\cos\beta\sin\beta + C_{yy}\sin^2\beta\right)\omega^4 \\ &+ \frac{1}{12}\left(C_{xy}\cos\beta\sin\beta + \rho\cos^2\beta\sin^2\beta\right)\omega^4 + \mathcal{O}(\omega^6).\end{aligned} \tag{46}$$

In the specific case when $C_{xy} = 0$, we can note that the choice $\rho = 0$ implies that

$$C_{xxyy} = \frac{C_{xx} + C_{yy}}{6}, \tag{47}$$

and we obtain a Fourier transform in which the fourth-order terms have a similar angular dependency, here an elliptic shape, as the second-order terms. Indeed, this elliptic shape also coincides with the shape of the corresponding Fourier transform of the Gaussian kernel on a continuous spatial domain.

In the specific case when $C_{xx} = C_{yy} = 1$ and $C_{xy} = 0$, the choice of $C_{xxyy} = (C_{xx} + C_{yy})/3$ does in turn correspond to the following discretization of the isotropic diffusion equation (Lindeberg [37, 39])

$$\partial_s L = \frac{1}{2}\left((1-\gamma)\,\nabla_5^2 L + \gamma\,\nabla_{\times 2}^2 L\right) \tag{48}$$

with

$$(\nabla_5^2 f)_{0,0} = f_{-1,0} + f_{+1,0} + f_{0,-1} + f_{0,+1} - 4f_{0,0} \tag{49}$$

$$(\nabla_{\times 2}^2 f)_{0,0} = 1/2(f_{-1,-1} + f_{-1,+1} + f_{+1,-1} + f_{+1,+1} - 4f_{0,0}), \tag{50}$$

for $\gamma = \frac{1}{3}$ and corresponding to the following composed computational molecule (Lindeberg [39, figure 4.5(d)])

$$\frac{2}{3}\nabla_5^2 + \frac{1}{3}\nabla_{\times 2}^2 = \frac{1}{6}\begin{pmatrix} 1 & 4 & 1 \\ 4 & -20 & 4 \\ 1 & 4 & 1 \end{pmatrix} \tag{51}$$

which gives the 3×3 discrete approximation of the Laplacian operator with highest degree of rotational symmetry.

4.5 Choice of the free parameter C_{xxyy}

In view of the above analysis, specifically with the condition (35) on a necessary relations between the free parameter C_{xxyy} in relation to the parameters C_{xx}, C_{xy} and C_{yy} of the covariance matrix, we can choose

$$C_{xxyy} = |C_{xy}|. \tag{52}$$

This is the minimal choice and leads to the discrete kernel with lowest fourth-order moment. (Note that for continuous Gaussian kernels, the fourth-order moment should be zero.)

The bottom row in figure 6 shows examples of discrete affine kernels generated in this way from the explicit expression of the Fourier transform (34) and with $C_{xxyy} = |C_{xy}|$ according to (52).

Alternatively, one can use the degrees of freedom in the free parameter C_{xxyy} to give a better approximation of rotational symmetry in the isotropic case or generalizations thereof in the affine case (47) and (46) or some other measure of the deviations from affine isotropy in combination with and restricted by the necessary positivity condition (35). Since this becomes technically more involved and also the degree of higher order affine approximation starts interfering with the positivity requirements already for rather low eccentricities $\lambda_{max}/\lambda_{min} > 2$, we do not develop that theory further in this treatment.

4.6 Discrete approximation of scale-space derivatives

Given the above discrete affine Gaussian kernels, we can in turn compute discrete approximations to directional derivatives according to

$$\delta_{\varphi} = \cos\varphi\, \delta_x + \sin\varphi\, \delta_y \tag{53}$$

$$\delta_{\perp\varphi} = -\sin\varphi\, \delta_x + \cos\varphi\, \delta_y \tag{54}$$

$$\delta_{\varphi\varphi} = \cos^2\varphi\, \delta_{xx} + 2\cos\varphi\, \sin\varphi\, \delta_{xy} + \sin^2\varphi\, \delta_{yy} \tag{55}$$

$$\delta_{\varphi\perp\varphi} = -\cos\varphi\, \sin\varphi\, \delta_{xx} - (\cos^2\varphi - \sin^2\varphi)\, \delta_{xy} + \cos\varphi\, \sin\varphi\, \delta_{yy} \tag{56}$$

$$\delta_{\perp\varphi\perp\varphi} = \sin^2\varphi\, \delta_{xx} - 2\cos\varphi\, \sin\varphi\, \delta_{xy} + \cos^2\varphi\, \delta_{yy} \tag{57}$$

where δ_{xx}, δ_{xy} and δ_{yy} denote discrete approximations to the partial derivative operators ∂_{xx}, ∂_{xy} and ∂_{yy}, respectively, and we can preferably choose the same discrete approximation operators as used for discretizing the continuous affine diffusion equation (2) into the semi-discrete affine diffusion equation (28) as well as in the generator of its solution (30) and (34).

The middle and the top rows in figure 6 show examples of discrete affine derivative approximation kernels generated in this way from discrete affine Gaussian kernels generated from the explicit expression of the Fourier transform (34) and with $C_{xxyy} = |C_{xy}|$ according to (52).

By applying corresponding discrete derivative approximation kernels to the colour-opponent channels u and v according to a colour-opponent representation (7), we obtain discrete affine Gaussian colour-opponent directional derivative approximations analogous to the colour-opponent receptive fields shown in figure 2.

Note that when we apply these affine Gaussian derivative approximation receptive fields in practice, we do never compute these kernels explicitly. Instead, we apply the

compact support 3×3 discrete directional derivative approximation kernels (53)–(57) directly to the output of the discrete affine Gaussian smoothing as obtained after the FFT implementation of discrete affine Gaussian kernels.

5 Discretization over scale s into 3×3 iteration kernels

By further discretizing the semi-discrete affine diffusion equation (28) with respect to the scale parameter s using Euler's forward method (Dahlquist *et al.* [11])

$$\partial_s L = \frac{L^{k+1} - L^k}{\Delta s}, \tag{58}$$

we obtain the discrete iteration formula

$$L^{k+1} = L^k + \Delta s \, \partial_s L = 1 + \Delta s \left(\frac{1}{2} \left(C_{xx} \, \delta_{xx} L + 2 C_{xy} \, \delta_{xy} L + C_{yy} \, \delta_{yy} L \right) + \frac{C_{xxyy}}{4} \, \delta_{xxyy} L \right) \tag{59}$$

with the following computational molecule for forward iteration by a 3×3 filter kernel:

$$k(C_{xx}, C_{xy}, C_{yy}, C_{xxyy}, \Delta_s) = \begin{pmatrix} \frac{1}{4}(-C_{xy} + C_{xxyy})\Delta s & \frac{1}{2}(C_{yy} - C_{xxyy})\Delta s & \frac{1}{4}(+C_{xy} + C_{xxyy})\Delta s \\ \frac{1}{2}(C_{xx} - C_{xxyy})\Delta s & 1 - (C_{xx} + C_{yy} - C_{xxyy})\Delta s & \frac{1}{2}(C_{xx} - C_{xxyy})\Delta s \\ \frac{1}{4}(+C_{xy} + C_{xxyy})\Delta s & \frac{1}{2}(C_{yy} - C_{xxyy})\Delta s & \frac{1}{4}(-C_{xy} + C_{xxyy})\Delta s \end{pmatrix}. \tag{60}$$

This kernel has spatial mean vector

$$m = \begin{pmatrix} m_x \\ m_y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{61}$$

and spatial covariance matrix

$$\Sigma = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy} & \Sigma_{yy} \end{pmatrix} = \begin{pmatrix} C_{xx} \Delta s & C_{xy} \Delta s \\ C_{xy} \Delta s & C_{yy} \Delta s \end{pmatrix}. \tag{62}$$

Thereby, in analogy with the semi-discrete affine Gaussian scale-space concept over a continuous scale parameter s, also the covariances matrices obtained from repeated forward iteration of the 3×3 kernel by K steps will be exactly equal to the covariance matrices of the corresponding continuous affine Gaussian kernels

$$\Sigma_{composed} = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy} & \Sigma_{yy} \end{pmatrix} = \begin{pmatrix} C_{xx} \, K \, \Delta s & C_{xy} \, K \, \Delta s \\ C_{xy} \, K \, \Delta s & C_{yy} \, K \, \Delta s \end{pmatrix}. \tag{63}$$

5.1 Choice of scale step Δs

In this section, we will analyse how large scale steps Δs can be taken with the discrete iteration kernel (60) while preserving scale-space properties or transfers thereof.

5.1.1 Normalization of the parameters of the covariance matrix

Since the scale step Δs interacts with the parameters C_{xx}, C_{xy}, C_{yy}, C_{xxyy} along a cone, we shall specifically choose to normalize the parameters C_{xx}, C_{xy}, C_{yy} and C_{xxyy} such that the maximum eigenvalue of the covariance matrix according to the parameterization (36), (37) and (38) is given by

$$\lambda_{max} = \max(\lambda_1, \lambda_2) = 1 \tag{64}$$

and thus the second eigenvalue is

$$\lambda_{min} \in [0, 1]. \tag{65}$$

With regard to our goal of discretizing the manifold of affine Gaussian receptive fields over all covariance matrices Σ as illustrated in figure 5, this parametrization has a very natural geometric interpretation. Assume that the central point in a first image f of a scene shows a local surface pattern viewed from a point along the direction of the surface normal at some depth Z, and that we compute isotropic Gaussian receptive field responses at scale s with the parameters of the covariance matrix given by $C_{xx} = C_{yy} = 1$ and $C_{xy} = 1$. Next, assume that we view the same local surface pattern from the same distance Z while from some other oblique direction with slant angle θ relative to the surface normal and corresponding to a local foreshortening factor of $\epsilon = \cos\theta$. Then, the second view should in an affine covariant manner be represented at scale s for a covariance matrix with eigenvalues $\lambda_{max} = 1$ and $\lambda_{min} \in [0, 1]$ with the parameter α that determines the eigendirections of the covariance matrix Σ is directly related to the local tilt direction.

In the following, we will throughout make use of such a parameterization of the covariance matrix Σ for expressing the magnitude of the scale step Δs.

5.1.2 Analysis for the case of the isotropic diffusion equation

To illustrate how the choice of the complementary filter parameter C_{xxyy} in (27) interrelates with how large scale steps one may take, let us initially consider a discretization of the isotropic diffusion equation for $C_{xx} = C_{yy} = 1$ and $C_{xy} = 0$ for the specific choice of $\gamma = 0$ in (48) and corresponding to $C_{xxyy} = 0$. This leads to a forward iteration kernel of the form

$$k(1, 0, 1, 0, \Delta s) = \begin{pmatrix} 0 & \frac{1}{2}\Delta s & 0 \\ \frac{1}{2}\Delta s & 1 - 2\Delta s & \frac{1}{2}\Delta s \\ 0 & \frac{1}{2}\Delta s & 0 \end{pmatrix}. \tag{66}$$

In (Lindeberg [37, 39]) a complete scale-space theory is developed for one-dimensional discrete signals and for isotropic higher-dimensional images over a symmetric spatial domain. In the one-dimensional version of this theory, spatial discretizations of the one-dimensional diffusion equation $\partial_s L = \frac{1}{2}\partial_{xx} L$ as obtained from Euler's forward method lead to kernels of the form

$$(\frac{\Delta s}{2}, 1 - \Delta s, \frac{\Delta s}{2}). \tag{67}$$

Based on the complete classification of discrete scale-space kernels over a one-dimensional domain, it is furthermore shown that such discretizations are true discrete scale-space

kernels in the one-dimensional sense of guaranteeing non-creation of new local extrema or new zero-crossings from finer to coarser temporal scales only if and only if

$$\Delta s \leq \frac{1}{2}. \tag{68}$$

Specifically, the boundary case $\Delta s = 1/2$ corresponds to forward iteration with the binomial kernel

$$\left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\right), \tag{69}$$

in which the central kernel value $1/2$ is twice as large as its nearest neighbours having values equal to $1/4$. If we apply a corresponding criterion to the two-dimensional forward iteration kernel (74), we obtain $1 - 2\Delta s \geq 2\Delta s/2$ leading to the requirement

$$\Delta s \leq \frac{1}{3}. \tag{70}$$

If we next again for the isotropic case of $C_{xx} = C_{yy} = 1$ and $C_{xy} = 0$, consider a discretization in the scale direction of (48) for a general value of the parameter $\gamma = C_{xxyy}$, we obtain a forward iteration kernel of the form (Lindeberg [39, Equation (4.39)])

$$k(1, 0, 1, \gamma, \Delta s) = \begin{pmatrix} \frac{1}{4}\gamma\Delta s & \frac{1}{2}(1-\gamma)\Delta s & \frac{1}{4}\gamma\Delta s \\ \frac{1}{2}(1-\gamma)\Delta s & 1-(2-\gamma)\Delta s & \frac{1}{2}(1-\gamma)\Delta s \\ \frac{1}{4}\gamma\Delta s & \frac{1}{2}(1-\gamma)\Delta s & \frac{1}{4}\gamma\Delta s \end{pmatrix}. \tag{71}$$

This kernel can be shown to be separable only if $\gamma = \Delta s$ (Lindeberg [39, Proposition 4.15]). The corresponding one-dimensional kernel $(a, 1-2a, a)$ for $a = \frac{1}{2}\sqrt{\gamma\Delta s}$ is then a one-dimensional discrete scale-space kernel only if

$$\Delta s \leq \frac{1}{2} \tag{72}$$

where the boundary case $\gamma = \Delta s = 1/2$ corresponds to forward iteration with the kernel

$$k(1, 0, 1, \tfrac{1}{2}, \tfrac{1}{2}) = \begin{pmatrix} \frac{1}{16} & \frac{1}{8} & \frac{1}{16} \\ \frac{1}{8} & \frac{1}{4} & \frac{1}{8} \\ \frac{1}{16} & \frac{1}{8} & \frac{1}{16} \end{pmatrix}. \tag{73}$$

This kernel is a common iteration kernel for computing pyramid representations (Crowley [8]; Crowley and Parker [10, 9]; Lindeberg [37, 39]; Lindeberg and Bretzner [48]). Specifically, applying the underlying one-dimensional kernel $(1/4, 1/2, 1/4)$ for separable convolution twice in the same direction gives a binomial kernel of the form $(1/16, 4/16, 6/16, 4/16, 1/16)$ within the family of kernels of length five derived by Burt and Adelson [6] $(\frac{1}{4} - \frac{a}{2}, \frac{1}{4}, a, \frac{1}{4}, \frac{1}{4} - \frac{a}{2})$ for $a = 3/8 = 0.375$ in good agreement with the numerical value $a \approx 0.4$ that they derived empirically.

Note that if we would apply the same value of $\Delta s = 1/2$ to the kernel (74) that arises from a spatial discretization of the isotropic diffusion equation for $\gamma = 0$ and corresponding to $C_{xxyy} = 0$, we would obtain a forward iteration kernel of the form

$$k(1, 0, 1, 0, \tfrac{1}{2}) = \begin{pmatrix} 0 & \frac{1}{4} & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & \frac{1}{4} & 0 \end{pmatrix}, \tag{74}$$

which would clearly not be a desirable kernel for scale-space filtering, for example, this kernel is not even unimodal. Thus, by the use of a non-zero value of γ and corresponding to a non-zero value of C_{xxyy}, in this specific case corresponding to $C_{xxyy} = C_{xx}/2 = C_{yy}/2$, we increase the value at the central position of the 3×3 kernel in relation to its nearest neighbours along the coordinate axes and are thereby able to take longer steps Δs in the scale direction.

5.1.3 Analysis for the non-isotropic diffusion equation

In this subsection, we consider discretizations of the affine Gaussian diffusion equation for general non-zero values of C_{xy} and C_{xxyy} and also without assuming $C_{xx} = C_{yy}$. Specifically, with a normalization of the parameters of the covariance matrix Σ to maximum eigenvalue equal, we will consider the possibility of taking scale steps up to $\Delta s = 1/2$ so that the discrete binomial kernel (73) that represents the maximum possible scale steps based on isotropic discrete scale-space theory is regarded as an upper bound on how much spatial smoothing is allowed in different spatial directions.

Requirement on the relations between the kernel values along the coordinate axes and the central point. If we inspired by the ratio of two between the kernel value $1/4$ at the center and the value $1/8$ at the nearest neighbouring points for the kernel (73) that represents the maximum possible smoothing for the isotropic diffusion equation, apply a corresponding criterion that the kernel values at the nearest four-neighbours to the central point having indices $(i, j) \in \{(-1, 0), (1, 0), (0, -1), (0, 1)\}$ should have values that are not greater than half the value of the central point with index $(0, 0)$, we obtain the conditions

$$\frac{1}{2}(C_{xx} - C_{xxyy})\,\Delta s \leq \frac{1}{2}(1 - (C_{xx}C_{yy} - C_{xxyy}))\,\Delta s, \tag{75}$$

$$\frac{1}{2}(C_{yy} - C_{xxyy})\,\Delta s \leq \frac{1}{2}(1 - (C_{xx}C_{yy} - C_{xxyy}))\,\Delta s, \tag{76}$$

which can be summarized into

$$(C_{xx} + C_{yy} + \max(C_{xx}, C_{yy}) - 2C_{xxyy})\,\Delta s \leq 1. \tag{77}$$

If we inspired by the above theoretical analysis for the isotropic case would like to be able to take scale steps up to $\Delta s = 1/2$ for general anisotropic discrete affine Gaussian kernels, we obtain the following condition for the values of the nearest four-neighbours to not exceed half the value at the central point

$$C_{xxyy} \geq \frac{C_{xx} + C_{yy} + \max(C_{xx}, C_{yy}) - 2}{2}, \tag{78}$$

where specifically the case $C_{xx} = C_{yy} = 1$ corresponds to $C_{xxyy} \geq 1/2$.

Requirement on the relations between the kernel values at the corners and the central point. If we apply a condition that the kernel values at the corners having indices $(i, j) \in \{(-1, -1), (-1, 1), (1, -1), (1, 1)\}$ should have values that are not greater than one fourth of the value at the central point, we obtain the condition

$$\frac{1}{4}(|C_{xy}| + C_{xxyy})\,\Delta s \leq \frac{1}{4}(1 - (C_{xx} + C_{yy} - C_{xxyy})\,\Delta s) \tag{79}$$

which can be simplified to to

$$(C_{xx} + C_{yy} + |C_{xy}|)\,\Delta s \leq 1. \tag{80}$$

If we again inspired by the theoretical analysis for the isotropic case would like to take scale steps up to $\Delta s = 1/2$ for general anisotropic discrete affine Gaussian kernels, then we obtain the following condition for the values at the corners should not exceed one quarter of the value at the central point

$$(C_{xx} + C_{yy} + |C_{xy}|) \leq 2. \tag{81}$$

For the specific assumption about the parameters C_{xx}, C_{xy} and C_{yy} corresponding to a covariance matrix with maximum eigenvalue $\lambda_{max} = 1$ and the other eigenvalue $\lambda_{min} \in [0, 1]$, this condition can according to (36), (37) and (38) be written

$$1 + \lambda_{min} + (1 - \lambda_{min}) \, |\sin \alpha \, \cos \alpha| \leq 2. \tag{82}$$

The worst case condition arises when $|\sin \alpha| = |\cos \alpha|$, for which the inequality can be reduced to

$$\lambda_{min} \leq 1 \tag{83}$$

and which is consistent with the assumptions thus showing that the condition (81) is guaranteed to hold for all relevant combinations of C_{xx}, C_{xy} and C_{yy}.

5.2 Choice of the free parameter C_{xxyy}

In view of the above analysis, specifically with the conditions (35) and (78) on necessary relations between the free parameter C_{xxyy} in relation to the parameters C_{xx}, C_{xy} and C_{yy} of the covariance matrix, we can given the a priori choice of scale step $\Delta s = 1/2$ choose

$$C_{xxyy} = \max \left(|C_{xy}|, \frac{C_{xx} + C_{yy} + \max(C_{xx}, C_{yy}) - 2}{2} \right). \tag{84}$$

This is the minimal choice and leads to the discrete kernel with lowest fourth-order moment. (Note again that for continuous Gaussian kernels, the fourth-order moment should be zero.)

The bottom row in figure 7 shows examples of discrete affine kernels generated in this way from repeated application of 3×3-kernels of the form (60) for $\Delta s = 1/2$ and with C_{xxyy} according to (84).

Alternatively, one can to choose the degrees of freedom within the constraint

$$\max \left(|C_{xy}|, \frac{C_{xx} + C_{yy} + \max(C_{xx}, C_{yy}) - 2}{2} \right) \leq C_{xxyy} \leq \min(C_{xx}, C_{yy}) \tag{85}$$

to minimize the deviations of the fourth-order terms in the Fourier transform (46) from an elliptic behaviour or some other measure of the deviations from affine isotropy. Since such an analysis does, however, become technically more complicated, we will not develop that theory further in this treatment.

5.3 Discrete approximation of scale-space derivatives

When computing discrete approximations to directional derivatives of these discrete affine Gaussian kernels, we can proceed in a similar way as for the discrete affine Gaussian kernels obtained via an FFT and define discrete directional derivative operators δ_φ, $\delta_{\perp\varphi}$, $\delta_{\varphi\varphi}$, $\delta_{\varphi\perp\varphi}$ and $\delta_{\perp\varphi\perp\varphi}$ according to equations (53)–(57).

The middle and the top rows in figure 7 show examples of discrete affine derivative approximation kernels generated in this way from repeated iteration of 3×3-kernels

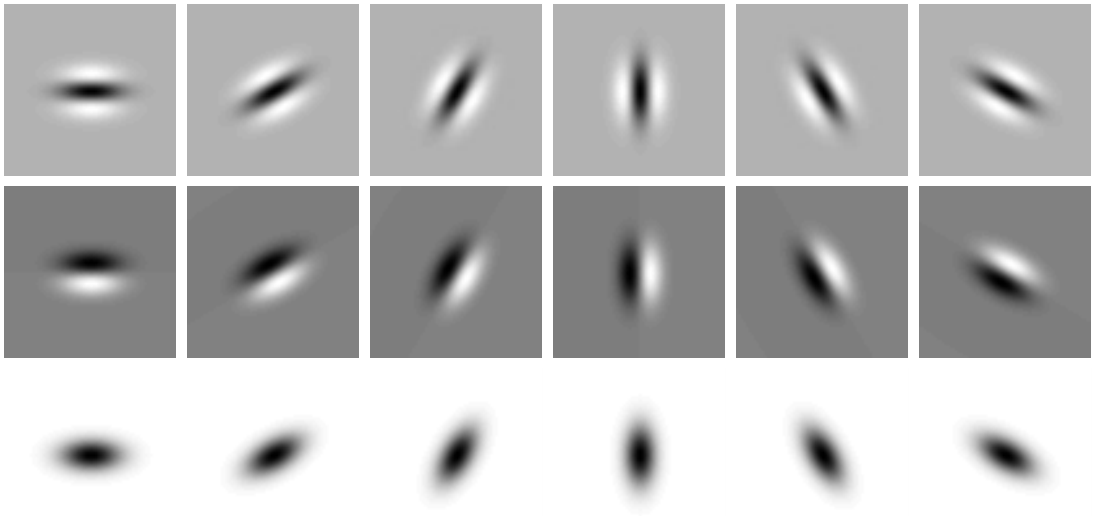

Figure 7: Examples of discrete affine Gaussian kernels $h(x, y;\ \Sigma)$ and their equivalent directional derivative approximation kernels $\delta_{\perp\varphi} h(x, y;\ \Sigma)$ and $\delta_{\perp\varphi\perp\varphi} h(x, y;\ \Sigma)$ up to order two in the two-dimensional case, here as generated from repeated iteration of 3×3-kernels of the form (60) for $\lambda_1 = 64$, $\lambda_2 = 16$ $\alpha = 0, \pi/6, \pi/3, \pi/2, 2\pi/3, 5\pi/6$, $\Delta s = 1/2$ and with C_{xxyy} according to (84). (Kernel size: 65×65 pixels.)

of the form (60) for $\lambda_1 = 64$, $\lambda_2 = 16$ $\alpha = 0, \pi/6, \pi/3, \pi/2, 2\pi/3, 5\pi/6$, $\Delta s = 1/2$ and with C_{xxyy} according to (84).

By applying corresponding discrete derivative approximation kernels to the colour-opponent channels u and v according to a colour-opponent representation (7), we obtain discrete affine Gaussian colour-opponent directional derivative approximations analogous to the colour-opponent receptive fields shown in figure 2.

Note that when we apply these affine Gaussian derivative approximation receptive fields in practice, we do never generate these kernels explicitly. Instead, we apply compact support 3×3 discrete directional derivative approximation kernels (53)–(57) directly to the output of iterative application of the 3×3-kernel (60), preferably combined with a spatial subsampling step at coarser levels of scales to implement an affine hybrid pyramid representation (as will be described in next section).

6 Affine hybrid pyramid

In this section, we will show how repeated filtering with the 3×3 kernel $k(C_{xx}, C_{xy}, C_{yy}, C_{xxyy}, \Delta_s)$ according to (60) can be combined with spatial subsampling operations to compute the discrete affine Gaussian scale-space representations at coarser scales in a computationally more efficient manner compared using the same resolution at all levels of scale. By applying the scale step $\Delta s \leq 1/2$ relative to the current resolution $h = 2^M$ after M spatial subsampling stages by a factor of 2, the effective bound on the spatial step will become $\Delta s \leq 2^{2M}/2$, thus implying larger scale steps and fewer iterations to reach coarser scale representations. This will lead to an extension of isotropic pyramid and hybrid pyramid representations (Burt and Adelson [6]; Crowley [8]; Crowley and Parker [10, 9]; Lindeberg and Bretzner [48]) to affine hybrid pyramid representations.

6.1 Reduction operators

Following Burt and Adelson [6]; Crowley [8]; Crowley and Parker [10, 9] and Lindeberg and Bretzner [48]), let us describe the the transformation between two adjacent scale levels in a pyramid by a reduction operator. For simplicity, let us assume that the pyramid is separable and that the size N of the smoothing filter is the same over both the coordinate directions and odd. Then, the transformation from the representation $L^{(i)}$ at the current scale level i, to the representation $L^{(i+1)}$ at the next coarser level $i+1$ is for some set of filter coefficients $c\colon \mathbb{Z} \to \mathbb{R}$ given by

$$L^{(i+1)} = \textsc{ReduceCycle}(L^{(i)}) \tag{86}$$

$$L^{(i+1)}(x) = \sum_{m=-(N-1)/2}^{(N-1)/2} \sum_{n=-(N-1)/2}^{(N-1)/2} c(m,n)\, L^{(i)}(sx-m, sy-n). \tag{87}$$

Next, let us assume that the smoothing operation can be decomposed into several smoothing steps:

$$\begin{array}{rcl} \textsc{ReduceCycle} & := & \textsc{SubSample} \\ & & \textsc{Smooth}^{+} \end{array} \tag{88}$$

where the notation $\textsc{Op}^{+}$ means that several operators of the form $\textsc{Op}$ may occur. $\textsc{ReduceCycle}$ is thus composed of one or more smoothing operations followed by a subsampling. The subsampling operation is here defined by

$$S = \textsc{SubSample}(L;\ s) \tag{89}$$

$$S(x,y) = L(sx, sy) \qquad (s \in \mathbb{Z}_{+}). \tag{90}$$

(where we usually choose $s = 2$) and each smoothing step according to

$$S = \textsc{Smooth}(L) \tag{91}$$

$$S(x,y) = \sum_{m=-N}^{N} \sum_{n=-N}^{N} c(m,n)\, L(x-m, y-n). \tag{92}$$

Specifically, we assume that the coefficients of the smoothing operation originate from the 3×3-kernel (60) arising from the discretization of the diffusion operator repeated K times

$$\textsc{Smooth}(L) = \textsc{DeltaSmooth}(L;\ \Delta s, K) = [\textsc{DeltaSmooth}(L;\ \Delta s, 1)]^{K}. \tag{93}$$

6.2 Equivalent convolution and derivative approximation kernels

Since the representation at each level is constructed from a set of repeated smoothing and subsampling operations, which are all linear operations, the composed operation can equivalently be modeled as the result of applying one kernel $\kappa^{(i,j)}$, termed *equivalent convolution kernel*, to the original image, followed by a pure subsampling step. If we define a dual operator to the $\textsc{ReduceCycle}$ operator according to

$$\begin{array}{rcl} \textsc{ExpandCycle} & := & \textsc{Smooth}^{+} \\ & & \textsc{Enlarge} \end{array}$$

where the $\textsc{Enlarge}$ operation enlarges any D-dimensional image by a factor 2

$$E = \textsc{Enlarge}(L) \tag{94}$$

$$E(x,y) = \begin{cases} L(x/2, y/2) & \text{if all indices in } x \text{ are multiples of 2} \\ 0 & \text{if any index in } x \text{ is not a multiple of 2} \end{cases} \tag{95}$$

the equivalent convolution kernel corresponding to level (i, j) can be written

$$\kappa^{(i,j)} = \textsc{ExpandAll}(\delta^{(i,j)}) \tag{96}$$

where $\delta^{(i,j)}$ is a discrete delta function at level (i, j) and ExpandAll denotes the ExpandCycle operators corresponding to the set of all the ReduceCycle operators used for reaching this level.

Similarly derivative approximations are computed by taking the grid spacing h at the current into explicit account

$$\partial_{x^r} \approx \mathcal{D}_{x^r} = \frac{1}{h^{|r|}} \, \delta_{x^r}, \tag{97}$$

at any level with resolution h in the pyramid, the corresponding *equivalent derivative approximation kernel* is given by

$$\kappa^{(i,j)}_{x^r} = \textsc{ExpandAll}(\delta^{(i,j)}_{x^r}) \tag{98}$$

where higher dimensional difference approximations $\delta_{x^r} = \delta_{x_1{}^{r_1}} \delta_{x_2{}^{r_2}} .. \delta_{\ x_D{}^{r_D}}$ are expressed in terms of the one-dimensional rth order difference operator

$$\delta_{x^r} = \begin{cases} (\delta_{xx})^{r/2} & \text{if } r \text{ is even} \\ \delta_x \, \delta_{x^{r-1}} & \text{if } r \text{ is odd.} \end{cases} \tag{99}$$

6.3 Measuring the subsampling rate

To describe how the grid spacing h depends on the scale parameter t in a hybrid pyramid, we introduce a *subsampling factor* ρ from the relation

$$h_{max} = \rho \, \min(\sigma_1, \sigma_2) = \rho \, \min(\sqrt{\lambda_1}, \sqrt{\lambda_2}) \tag{100}$$

where for reasons of computational efficiency we define the actual grid spacing as the maximum power of two that does not exceed this upper bound

$$h(t, \rho) = \begin{cases} \max_{h'=2^{i-1} \colon i \in \mathbb{Z}_+ \setminus \{0\}} h' : h' < h_{max}(\lambda_1, \lambda_2, \rho) & \text{if } h_{max} \geq 1, \\ 1 & \text{otherwise.} \end{cases} \tag{101}$$

Thus, a subsampling factor of $\rho = 0$ corresponds to preserving the original resolution at all levels of scales, whereas increasing values of ρ correspond to higher degrees of subsampling at coarser levels of scale.

7 Scale-normalized derivatives for affine Gaussian scale space

Continuous affine Gaussian derivative kernels. When defining scale-normalized derivatives for the continuous affine Gaussian receptive fields of the form $g_{\varphi^m \perp \varphi^n}(x, y; \ \Sigma_s)$, we use the eigenvalues of the spatial covariance matrix Σ_s as scale parameters

$$L_{\varphi^m \perp \varphi^n, norm}(x, y; \ \Sigma_s) = \lambda_1^{m\gamma_1} \, \lambda_2^{n\gamma_2} \, \partial^m_{\varphi} \partial^n_{\perp \varphi} \, L(x, y; \ \Sigma_s), \tag{102}$$

where γ_1 and γ_2 denote possibly different scale normalization powers for the two orthogonal directions φ and $\perp\varphi$ and specifically the choice $\gamma_1 = \gamma_2 = 1$ implies maximal scale invariance.

Discrete affine Gaussian derivative kernels. For discrete approximations of affine Gaussian derivatives obtained by applying the discrete directional derivative approximations according to (53)–(57) to either the semi-discrete affine scale-space concept according to (28) or the scale-discretized scale-space concept corresponding to repeated application of the 3×3-kernel (60), we can in a corresponding manner define *variance-normalized* affine Gaussian scale-space derivatives according to

$$L_{\varphi^m \perp \varphi^n, norm}(x, y;\ \Sigma_s) = \lambda_1^{m\gamma_1}\, \lambda_2^{n\gamma_2}\, \delta_\varphi^m\ \delta_\varphi^n\, L(x, y;\ \Sigma_s), \tag{103}$$

with the implicit understanding that $\delta_\varphi^2 = \delta_{\varphi\varphi}$ and $\delta_{\perp\varphi}^2 = \delta_{\perp\varphi\perp\varphi}$.

Alternatively, we can define scale-normalized affine Gaussian derivatives from l_p-normalization by requiring that the discrete l_p-norm of the discrete Gaussian derivative kernel being equal to the continuous L_p-norm of the corresponding continuous Gaussian derivative operator

$$L_{\varphi^m \perp \varphi^n, norm}(x, y;\ \Sigma_s) = \alpha(\lambda_1, \gamma_1, \lambda_2, \gamma_2)\, \delta_\varphi^m\ \delta_\varphi^n\, L(x, y;\ \Sigma_s), \tag{104}$$

with $\alpha(\lambda_1, \gamma_1, \lambda_2, \gamma_2)$ determined such that

$$\alpha(\lambda_1, \gamma_1, \lambda_2, \gamma_2)\, \|\delta_\varphi^m\ \delta_\varphi^n\, h(x, y;\ \Sigma_s)\|_1 = \lambda_1^{m\gamma_1}\, \lambda_2^{n\gamma_2}\, \|\partial_\varphi^m \partial_{\perp\varphi}^n\, g(x, y;\ \Sigma_s)\|_1. \tag{105}$$

For scale-normalized derivatives in a hybrid pyramid representation, the definitions are analogous with only difference that the expression $\|\delta_\varphi^m\ \delta_\varphi^n\, h(x, y;\ \Sigma_s)\|_1$ in (105) should be replaced by $\|\kappa_{x^r}^{(i,j)}\|_1$ according to (98).

8 Summary and discussion

We have presented a theory for discretizing the affine Gaussian scale-space concept so that scale-space properties hold also in a discrete sense for the discrete implementation. Specifically, we have presented two types of spatial discretizations of the continuous affine diffusion equation that describes the effect of convolution with affine Gaussian kernels: (i) a semi-discretization over space only while leaving the continuous scale parameter in terms of a spatial covariance matrix defined over a continuum and (ii) an additional discretization in the scale direction that leads to compact 3×3 kernels for repeated application and with their parameter having immediate interpretations in terms of the spatial covariance matrix of the composed smoothing effect and which optionally can be combined with spatial subsampling operations to compute affine hybrid pyramid representations. These two families of discrete affine Gaussian kernels can in turn be combined with discretizations of spatial directional derivative operators for computing discrete approximations of directional derivatives, leading to families of discrete affine Gaussian derivatives which are intended to serve as a basis for expressing later stage visual processes.

A limitation of the presented genuinely discrete affine Gaussian scale-space concepts is that the requirement of a positive discretization as implied by the discrete scale-space axioms underlying the formulation of the theory implies a bound on the ratio between the maximum and minimum eigenvalues of the spatial covariance matrix that determines the shape of the affine Gaussian kernels. By theoretical analysis we have shown that this upper bound $\epsilon = \lambda_{max}/\lambda_{min} = 3 + 2\sqrt{2} \approx 5.8$ corresponds to the amount of foreshortening caused by a slant angle of 65.5 degrees, which is within the bounds of basic use cases for expressing affine covariant image operations between multiple views of the same object.

If stronger affine image deformations need to be handled by the vision system, then alternative ways of discretizing the theory should be seeked. In our earlier work on affine shape adaptation, we have in such situations explored the affine covariant property of the continuous affine Gaussian scale-space concept (12) by first warping the original image by an affine transformation implemented in terms of spline interpolation, performing separable discrete scale-space smoothing in the warped domain and then warping the smoothed image back to the original image domain. Such an operation is usually faster than first generating a filter corresponding to the sampled affine Gaussian kernel or better using a locally integrated affine Gaussian kernel over the support region of each pixel and then performing inseparable convolution with the resulting discrete kernel. This form of discretization has also been used in most implementations of affine shape adaptation for interest point detection (Lindeberg and Gårding [50]; Mikolajczyk and Schmid [52]; Tuytelaars and Mikolajczyk [67] Lindeberg [43]). Another alternative proposed by Geusebroek *et al.* [20] is to approximate the inseparable affine Gaussian kernel with a set of separable recursive filters. For implementations of these types, discrete scale-space properties are, however, not guaranteed to hold, so the relation to the underlying aims of scale-space theory rests more on the degree of accuracy by which equations obtained from the continuous scale-space theory are numerically approximated.

A Non-isotropic scale-space for discrete signals: Necessity and sufficiency

This appendix presents a general theoretical result concerning linear shift-invariant scale-space representations of discrete signals; the result is a generalization of earlier results presented in (Lindeberg [37, 39]) for isotropic discrete scale space.

A.1 Definitions

Let us summarize this (minimal) set of basic properties a family should satisfy to be a candidate family for generating a linear scale-space.

Definition 1 *(Discrete pre-scale-space family of kernels)*
A one-parameter family of kernels $T\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ *is said to be a discrete pre-scale-space family of kernels if it satisfies*

- $T(\cdot;\ 0) = \delta(\cdot)$,
- *the semi-group property* $T(\cdot;\ s_1) * T(\cdot;\ s_2) = T(\cdot;\ s_1 + s_2)$,
- *the continuity requirement* $\| T(\cdot;\ s) - \delta(\cdot) \|_1 \to 0$ *when* $s \downarrow 0$.

Definition 2 *(Discrete pre-scale-space representation) Let* $f\colon \mathbb{Z}^N \to \mathbb{R}$ *be a discrete signal and let* $T\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ *be a discrete pre-scale-space family of kernels. Then, the one-parameter family of signals* $L\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ *given by*

$$L(x;\ s) = \sum_{\xi \in \mathbb{Z}^N} T(\xi;\ s)\, f(x - \xi). \tag{106}$$

is said to be the discrete pre-scale-space representation of f *generated by* T.

Lemma 3 *(A discrete pre-scale-space representation is differentiable)*
Let $L\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ be the discrete pre-scale-space representation of a signal $f\colon \mathbb{Z}^N \to \mathbb{R}$ in l_1. Then, L satisfies the differential equation

$$\partial_s L = \mathcal{A} L \tag{107}$$

for some linear and shift-invariant operator $\mathcal{A}$.

Proof: If f is sufficiently regular, *e.g.*, if $f \in L_1$, define a family of operators $\{\mathcal{T}_s, s > 0\}$, here from from L_1 to L_1, by $\mathcal{T}_s f = T(\cdot;\ s) * f$. Due to the conditions imposed on the kernels, the family satisfies the relation

$$\lim_{s \to s_0} \| (\mathcal{T}_s - \mathcal{T}_{s_0}) f \|_1 = \lim_{s \to s_0} \| (\mathcal{T}_{s-s_0} - \mathcal{I})(\mathcal{T}_{s_0} f) \|_1 = 0, \tag{108}$$

where $\mathcal{I}$ is the identity operator. Such a family is called a strongly continuous semi-group of operators (Hille and Phillips [24, pages 58–59]) A semi-group is often characterized by its *infinitesimal generator* $\mathcal{A}$ defined by

$$\mathcal{A} f = \lim_{h \downarrow 0} \frac{\mathcal{T}_h f - f}{h}. \tag{109}$$

The set of elements f for which $\mathcal{A}$ exists is denoted $\mathcal{D}(\mathcal{A})$. This set is not empty and never reduces to the zero element. Actually, it is even dense in L_1 (Hille and Phillips [24, page 307]). If this operator exists then

$$\begin{aligned} \lim_{h \downarrow 0} \frac{L(\cdot, \cdot;\ s+h) - L(\cdot, \cdot;\ s)}{h} &= \lim_{h \downarrow 0} \frac{\mathcal{T}_{s+h} f - \mathcal{T}_s f}{h} \\ &= \lim_{h \downarrow 0} \frac{\mathcal{T}_h(\mathcal{T}_s f) - (\mathcal{T}_s f)}{h} = \mathcal{A}(\mathcal{T}_s f) = \mathcal{A} L(\cdot;\ s). \end{aligned} \tag{110}$$

According to a theorem in Hille and Phillips [24, page 308] strong continuity implies $\partial_s(\mathcal{T}_s f) = \mathcal{A}\mathcal{T}_s f = \mathcal{T}_s \mathcal{A} f$ for all $f \in \mathcal{D}(\mathcal{A})$. Hence, the scale-space family L must obey the differential equation $\partial_s L = \mathcal{A} L$ for some linear operator $\mathcal{A}$. Since L is generated from f by a convolution operation it follows that $\mathcal{A}$ must be shift-invariant. □

Definition 4 *(Pre-scale-space property: Non-enhancement of local extrema)*
A discrete pre-scale-space representation $L\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ of a smooth (infinitely continuously differentiable) signal is said to possess pre-scale-space properties, or equivalently not to enhance local extrema, if for every value of the scale parameter $s_0 \in \mathbb{R}_+$ it holds that if $x_0 \in \mathbb{Z}^N$ is a critical point for the mapping $x \mapsto L(x;\ s_0)$ and if the Hessian matrix at this point is non-degenerate, then the derivative of L with respect to s satisfies

$$\partial_s L \leq 0 \qquad \text{at a non-degenerate local maximum,} \tag{111}$$

$$\partial_s L \geq 0 \qquad \text{at a non-degenerate local minimum.} \tag{112}$$

Definition 5 *(Discrete scale-space family of kernels) A one-parameter family of discrete pre-scale-space kernels $T\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ is said to be a discrete scale-space family of kernels if for any discrete signal $f\colon \mathbb{Z}^N \to \mathbb{R} \in l_1$ the discrete pre-scale-space representation of f generated by T possesses pre-scale-space properties,* i.e.*, if for any signal local extrema are never enhanced.*

Definition 6 *(Discrete scale-space representation) A discrete pre-scale-space representation $L\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ of a signal $f\colon \mathbb{Z}^N \to \mathbb{R}$ generated by a family of discrete kernels $T\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$, which are discrete scale-space kernels, is said to be a discrete scale-space representation of f.*

A.2 Necessity

Theorem 7 *(Scale-space for discrete signals: Necessity)*
A discrete scale-space representation $L\colon \mathbb{Z}^N \times \mathbb{R}_+ \to \mathbb{R}$ *of a discrete signal* $f\colon \mathbb{Z}^N \to \mathbb{R}$ *satisfies the differential equation*

$$\partial_s L = \mathcal{A} L \tag{113}$$

with initial condition $L(\cdot;\ 0) = f(\cdot)$ *for some infinitesimal scale-space generator* $\mathcal{A}$ *characterized by:*

- *the* locality *condition* $a_\xi = 0$ *if* $|\xi|_\infty > 1$,
- *the* positivity *constraint* $a_\xi \geq 0$ *if* $\xi \neq 0$, *and*
- *the* zero sum *condition* $\sum_{\xi \in \mathbb{Z}^D} a_\xi = 0$.

Proof: The proof consists of two parts. The first part has already been presented in Lemma 3, where it was shown that the requirements on the kernels imply that the family L obeys a linear differential equation. Because of the shift invariance $\mathcal{A}L$ can be written in the form

$$(\partial_s L)(x;\ s) = (\mathcal{A}L)(x;\ s) = \sum_{\xi \in \mathbb{Z}^D} a_\xi L(x + \xi;\ s), \tag{114}$$

In the second part counterexamples will be constructed from various simple test functions in order to delimit the class of possible operators.

D.1. The extremum point conditions (111), (112) combined with definitions 5-6 mean that $\mathcal{A}$ must be *local*, i.e., that $a_\xi = 0$ if $\xi \notin N_+(0)$. This is easily understood by studying the following counterexample: First, assume that $a_{\xi_0} > 0$ for some $\xi_0 \notin N_+(0)$ and define a function $f_1\colon \mathbb{Z}^N \to \mathbb{R}$ by

$$f_1(x) = \begin{cases} \varepsilon > 0 & \text{if } x = 0, \\ 0 & \text{if } x \in N(\xi_0), \\ 1 & \text{if } x = \xi_0, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \tag{115}$$

Obviously, ξ_0 is a local maximum point for f_1. From (13) one obtains $\partial_s L(\xi_0;\ 0) = \epsilon a_0 + a_{\xi_0}$. It is clear that this value can be positive provided that ε is chosen small enough. Hence, L cannot satisfy (111). Similarly, it can also be shown that $a_{\xi_0} < 0$ leads to a violation of the non-enhancement property (112) (let $\varepsilon < 0$). Consequently, a_ξ must be zero if $\xi \notin N_+(0)$.

D.2. With the given (weak) definitions of local extremum points it is clear that 0 is both a local maximum point and a local minimum point. Hence $\partial_s L(0;\ 0)$ must be zero, and the *coefficients sum to zero*

$$\sum_{\xi \in \mathbb{Z}^N} a_\xi = 0, \tag{116}$$

D.3. Finally, by observing that due to the zero sum condition, (13) can be written

$$\partial_s L = (\mathcal{A}L)(x;\ s) = \sum_{\xi \in N(0)} a_\xi (L(x - \xi;\ s) - L(x;\ s)), \tag{117}$$

and by considering the test function

$$f_3(x, y) = \begin{cases} \epsilon > 0 & \text{if } x = 0, \\ -1 & \text{if } x = \tilde{\xi}, \text{ and} \\ 0 & \text{otherwise,} \end{cases} \tag{118}$$

for some $\tilde{\xi}$ in $N(0)$ one easily realizes that a_ξ must be *non-negative* if $\xi \in N(0)$. The initial condition $L(\cdot;\ 0) = f$ is a direct consequence of the definition of pre-scale-space kernel. □

A.3 Sufficiency

Theorem 8 *(Scale-space for discrete signals: Sufficiency)*
Given an infinitesimal generator $\mathcal{A}$ corresponding to

$$(\mathcal{A}L)(x;\ s) = \sum_{\xi \in \mathbb{Z}^D} a_\xi L(x + \xi;\ s), \tag{119}$$

where the coefficients a_ξ satisfy

- *the* locality *condition $a_\xi = 0$ if $|\xi|_\infty > 1$,*
- *the* positivity *constraint $a_\xi \geq 0$ if $\xi \neq 0$, and*
- *the* zero sum *condition $\sum_{\xi \in \mathbb{Z}^D} a_\xi = 0$,*

the solution of the diffusion equation

$$\partial_s L = \mathcal{A}L \tag{120}$$

with initial condition $L(\cdot;\ 0) = f$ constitutes a discrete scale-space representation. Specifically, L obeys

$$\partial_s L \leq 0 \qquad \text{at a non-degenerate local maximum,} \tag{121}$$

$$\partial_s L \geq 0 \qquad \text{at a non-degenerate local minimum.} \tag{122}$$

Proof: From (117) it follows that the influence of L at any extremum point $(x_0;\ s_0)$ in scale-space can be written

$$\partial_s L(x_0;\ s_0) = (\mathcal{A}L)(x_0;\ s_0) = \sum_{\xi \in N(0)} a_\xi (L(x_0 - \xi;\ s_0) - L(x_0;\ s_0)), \tag{123}$$

If x_0 is a local maximum, then all differences $L(x_0 - \xi;\ s_0) - L(x_0;\ s_0)$ are negative. Combined with the positivity of a_ξ it follows that $\partial_s L(x_0;\ s_0) < 0$. If x_0 is a local minimum, we apply the same way of reasoning to $-L$. □

References

[1] Almansa, A., Lindeberg, T.: Fingerprint enhancement by shape adaptation of scale-space operators with automatic scale-selection. IEEE Trans. on Image Processing **9**(12), 2027–2042 (2000)

[2] Baumberg, A.: Reliable feature matching across widely separated views. In: Proc. Computer Vision and Pattern Recognition (CVPR'00), pp. I:1774–1781. Hilton Head, SC (2000)

[3] Bay, H., Ess, A., Tuytelaars, T., van Gool, L.: Speeded up robust features (SURF). Computer Vision and Image Understanding **110**(3), 346–359 (2008)

[4] Bouma, H., Vilanova, A., Bescós, J.O., ter Haar Romeny, B., Gerritsen, F.A.: Fast and accurate Gaussian derivatives based on B-splines. In: International Conference on Scale Space and Variational Methods in Computer Vision (SSVM 2007), pp. 406–417 (2007). Springer LNCS volume 4485

[5] Burghouts, G.J., Geusebroek, J.M.: Performance evaluation of local colour invariants. Computer Vision and Image Understanding **113**(1), 48–62 (2009)

[6] Burt, P.J., Adelson, E.H.: The Laplacian pyramid as a compact image code. IEEE Trans. Communications **9:4**, 532–540 (1983)

[7] Conway, B.R., Livingstone, M.S.: Spatial and temporal properties of cone signals in alert macaque primary visual cortex. Journal of Neuroscience **26**(42), 10,826–10,846 (2006)

[8] Crowley, J.L.: A representation for visual information. Ph.D. thesis, Carnegie-Mellon University, Robotics Institute, Pittsburgh, Pennsylvania (1981)

[9] Crowley, J.L., Parker, A.C.: A representation for shape based on peaks and ridges in the Difference of Low-Pass Transform. IEEE Trans. Pattern Analysis and Machine Intell. **6**(2), 156–170 (1984)

[10] Crowley, J.L., Stern, R.M.: Fast computation of the Difference of Low Pass Transform. IEEE Trans. Pattern Analysis and Machine Intell. **6**, 212–222 (1984)

[11] Dahlquist, G., Björk, Å., Anderson, N.: Numerical Methods. Prentice-Hall (1974)

[12] Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Proc. Computer Vision and Pattern Recognition (CVPR 2005), vol. 1, pp. 886–893 (2005)

[13] DeAngelis, G.C., Anzai, A.: A modern view of the classical receptive field: Linear and non-linear spatio-temporal processing by V1 neurons. In: L.M. Chalupa, J.S. Werner (eds.) The Visual Neurosciences, vol. 1, pp. 704–719. MIT Press (2004)

[14] DeAngelis, G.C., Ohzawa, I., Freeman, R.D.: Receptive field dynamics in the central visual pathways. Trends in Neuroscience **18**(10), 451–457 (1995)

[15] Deriche, R.: Using Canny's criteria to derive a recursively implemented optimal edge detector. International Journal of Computer Vision **1**, 167–187 (1987)

[16] Dollár, P., Appel, R., Belongie, S., Perona, P.: Fast feature pyramids for object detection. IEEE Transactions on Pattern Analysis and Machine Intelligence **36**(8), 1532–1545 (2014)

[17] Farid, H., Simoncelli, E.P.: Differentiation of discrete multidimensional signals. IEEE Transactions on Image Processing **13**(4), 496–508 (2004)

[18] Florack, L.: A spatio-frequency trade-off scale for scale-space filtering. IEEE Transactions on Pattern Analysis and Machine Intelligence **22**(9), 1050–1055 (2000)

[19] Florack, L.M.J.: Image Structure. Series in Mathematical Imaging and Vision. Springer (1997)

[20] Geusebroek, J.M., Smeulders, A.W.M., van de Weijer, J.: Fast anisotropic Gauss filtering. IEEE Transactions on Image Processing **12**(8), 938–943 (2003)

[21] ter Haar Romeny, B. (ed.): Geometry-Driven Diffusion in Computer Vision. Series in Mathematical Imaging and Vision. Springer (1994)

[22] Hall, D., de Verdiere, V., Crowley, J.: Object recognition using coloured receptive fields. In: Proc. European Conf. on Computer Vision (ECCV 2000), *Springer Lecture Notes in Computer Science*, vol. 1842, pp. I:164–177. Springer, Dublin, Ireland (2000)

[23] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. arXiv preprint arXiv:1512.03385 (2015)

[24] Hille, E., Phillips, R.S.: Functional Analysis and Semi-Groups, vol. XXXI. American Mathematical Society Colloquium Publications (1957)

[25] Hubel, D.H., Wiesel, T.N.: Receptive fields of single neurones in the cat's striate cortex. J Physiol **147**, 226–238 (1959)

[26] Hubel, D.H., Wiesel, T.N.: Brain and Visual Perception: The Story of a 25-Year Collaboration. Oxford University Press (2005)

[27] Iijima, T.: Observation theory of two-dimensional visual patterns. Tech. rep., Papers of Technical Group on Automata and Automatic Control, IECE, Japan (1962)

[28] Johnson, E.N., Hawken, M.J., Shapley, R.: The orientation selectivity of color-responsive neurons in Macaque V1. The Journal of Neuroscience **28**(32), 8096–8106 (2008)

[29] Koenderink, J.J.: The structure of images. Biological Cybernetics **50**, 363–370 (1984)

[30] Koenderink, J.J., van Doorn, A.J.: Representation of local geometry in the visual system. Biological Cybernetics **55**, 367–375 (1987)

[31] Koenderink, J.J., van Doorn, A.J.: Generic neighborhood operators. IEEE Trans. Pattern Analysis and Machine Intell. **14**(6), 597–605 (1992)

[32] Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)

[33] Larsen, A.B.L., Darkner, S., Dahl, A.L., Pedersen, K.S.: Jet-based local image descriptors. In: Proc. European Conference on Computer Vision (ECCV 2012), *Springer Lecture Notes in Computer Science*, vol. 7574, pp. III:638–650. Springer (2012)

[34] Lazebnik, S., Schmid, C., Ponce, J.: A sparse texture representation using local affine regions. IEEE Trans. Pattern Analysis and Machine Intell. **27**(8), 1265–1278 (2005)

[35] Lim, J.Y., Stiehl, H.S.: A generalized discrete scale-space formulation for 2-D and 3-D signals. In: International Conference on Scale-Space Theories in Computer Vision (Scale-Space'03), pp. 132–147 (2003). Springer LNCS volume 2695

[36] Linde, O., Lindeberg, T.: Composed complex-cue histograms: An investigation of the information content in receptive field based image descriptors for object recognition. Computer Vision and Image Understanding **116**, 538–560 (2012)

[37] Lindeberg, T.: Scale-space for discrete signals. IEEE Trans. Pattern Analysis and Machine Intell. **12**(3), 234–254 (1990)

[38] Lindeberg, T.: Discrete derivative approximations with scale-space properties: A basis for low-level feature extraction. Journal of Mathematical Imaging and Vision **3**(4), 349–376 (1993)

[39] Lindeberg, T.: Scale-Space Theory in Computer Vision. Springer (1993)

[40] Lindeberg, T.: Scale-space theory: A basic tool for analysing structures at different scales. Journal of Applied Statistics **21**(2), 225–270 (1994). Also available from http://www.csc.kth.se/∼tony/abstracts/Lin94-SI-abstract.html

[41] Lindeberg, T.: Linear spatio-temporal scale-space. In: B.M. ter Haar Romeny, L.M.J. Florack, J.J. Koenderink, M.A. Viergever (eds.) Proc. International Conference on Scale-Space Theory in Computer Vision (Scale-Space'97), *Springer Lecture Notes in Computer Science*, vol. 1252, pp. 113–127. Springer, Utrecht, The Netherlands (1997)

[42] Lindeberg, T.: Linear spatio-temporal scale-space. Tech. Rep. ISRN KTH/NA/P--01/22--SE, Dept. of Numerical Analysis and Computer Science, KTH (2001). Available from http://www.csc.kth.se/cvap/abstracts/cvap257.html

[43] Lindeberg, T.: Scale-space. In: B. Wah (ed.) Encyclopedia of Computer Science and Engineering, pp. 2495–2504. John Wiley and Sons, Hoboken, New Jersey (2008)

[44] Lindeberg, T.: Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space. Journal of Mathematical Imaging and Vision **40**(1), 36–81 (2011)

[45] Lindeberg, T.: A computational theory of visual receptive fields. Biological Cybernetics **107**(6), 589–635 (2013)

[46] Lindeberg, T.: Generalized axiomatic scale-space theory. In: P. Hawkes (ed.) Advances in Imaging and Electron Physics, vol. 178, pp. 1–96. Elsevier (2013)

[47] Lindeberg, T.: Time-causal and time-recursive spatio-temporal receptive fields. Journal of Mathematical Imaging and Vision **55**(1), 50–88 (2016)

[48] Lindeberg, T., Bretzner, L.: Real-time scale selection in hybrid multi-scale representations. In: L. Griffin, M. Lillholm (eds.) Proc. Scale-Space Methods in Computer Vision (Scale-Space'03), *Springer Lecture Notes in Computer Science*, vol. 2695, pp. 148–163. Springer, Isle of Skye, Scotland (2003)

[49] Lindeberg, T., Gårding, J.: Shape-adapted smoothing in estimation of 3-D depth cues from affine distortions of local 2-D structure. In: J.O. Eklundh (ed.) Proc. European Conf. on Computer Vision (ECCV'94), *Lecture Notes in Computer Science*, vol. 800, pp. 389–400. Springer-Verlag, Stockholm, Sweden (1994)

[50] Lindeberg, T., Gårding, J.: Shape-adapted smoothing in estimation of 3-D depth cues from affine distortions of local 2-D structure. Image and Vision Computing **15**, 415–434 (1997)

[51] Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2), 91–110 (2004)

[52] Mikolajczyk, K., Schmid, C.: Scale and affine invariant interest point detectors. International Journal of Computer Vision **60**(1), 63–86 (2004)

[53] Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., Kadir, T., van Gool, L.: A comparison of affine region detectors. International Journal of Computer Vision **65**(1–2), 43–72 (2005)

[54] Morel, J.M., Yu, G.: ASIFT: A new framework for fully affine invariant image comparison. SIAM Journal on Imaging Sciences **2**(2), 438–469 (2009)

[55] Rothganger, F., Lazebnik, S., Schmid, C., Ponce, J.: 3D object modeling and recognition using local affine-invariant image descriptors and multi-view spatial constraints. International Journal of Computer Vision **66**(3), 231–259 (2006)

[56] Sadek, R., Constantinopoulos, C., Meinhardt, E., C. Ballester, C., Caselles, V.: On affine invariant descriptors related to SIFT. SIAM Journal on Imaging Sciences **5**(2), 652–687 (2012)

[57] van de Sande, K.E.A., Gevers, T., Snoek, C.G.M.: Evaluating color descriptors for object and scene recognition. IEEE Trans. Pattern Analysis and Machine Intell. **32**(9), 1582–1596 (2010)

[58] Schiele, B., Crowley, J.: Recognition without correspondence using multidimensional receptive field histograms. International Journal of Computer Vision **36**(1), 31–50 (2000)

[59] Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., Poggio, T.: Robust object recognition with cortex-like mechanisms. IEEE Transactions on Pattern Analysis and Machine Intelligence **29**(3), 411–426 (2007)

[60] Shapley, R., Hawken, M.J.: Color in the cortex: single-and double-opponent cells. Vision research **51**(7), 701–717 (2011)

[61] Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (ICLR 2015) (2015). ArXiv preprint arXiv:1409.1556

[62] Slavík, A., Stehlík, P.: Dynamic diffusion-type equations on discrete-space domains. Journal of Mathematical Analysis and Applications **427**(1), 525–545 (2015)

[63] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proc. Computer Vision and Pattern Recognition (CVPR 2015), pp. 1–9 (2015)

[64] Tola, E., Lepetit, V., Fua, P.: Daisy: An efficient dense descriptor applied to wide baseline stereo. IEEE Trans. Pattern Analysis and Machine Intell. **32**(5), 815–830 (2010)

[65] Tschirsich, M., Kuijper, A.: Notes on discrete Gaussian scale space. Journal of Mathematical Imaging and Vision **51**, 106–123 (2015)

[66] Tuytelaars, T., van Gool, L.: Matching widely separated views based on affine invariant regions. International Journal of Computer Vision **59**(1), 61–85 (2004)

[67] Tuytelaars, T., Mikolajczyk, K.: A Survey on Local Invariant Features, *Foundations and Trends in Computer Graphics and Vision*, vol. 3(3). Now Publishers (2008)

[68] van Vliet, L.J., Young, I.T., Verbeek, P.W.: Recursive Gaussian derivative filters. In: International Conference on Pattern Recognition, vol. 1, pp. 509–514 (1998)

[69] Wang, Y.P.: Image representations using multiscale differential operators. IEEE Transactions on Image Processing **8**(12), 1757–1771 (1999)

[70] Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner-Verlag, Stuttgart, Germany (1998)

[71] Weickert, J.: Coherence-enhancing diffusion filtering. International Journal of Computer Vision **31**(2-3), 111–127 (1999)

[72] Weickert, J., ter Haar Romeny, B.M., Viergever, M.A.: Efficient and reliable schemes for nonlinear diffusion filtering. IEEE Transactions on Image Processing **7**(3), 398–410 (1998)

[73] Weickert, J., Ishikawa, S., Imiya, A.: Linear scale-space has first been proposed in Japan. Journal of Mathematical Imaging and Vision **10**(3), 237–252 (1999)

[74] Witkin, A.P.: Scale-space filtering. In: Proc. 8th Int. Joint Conf. Art. Intell., pp. 1019–1022. Karlsruhe, Germany (1983)

[75] Young, R.A.: The Gaussian derivative model for spatial vision: I. Retinal mechanisms. Spatial Vision **2**, 273–293 (1987)

[76] Young, R.A., Lesperance, R.M., Meyer, W.W.: The Gaussian derivative model for spatio-temporal vision: I. Cortical model. Spatial Vision **14**(3, 4), 261–319 (2001)